

09/234,427

Request to Enter Supplemental Amendment filed September 9, 2002

EXHIBIT A

BEST AVAILABLE COPY

Technical drawing of a mechanical part, likely a pump housing or valve body, showing two views: a top view and a side view. The top view is a rectangle with rounded corners and a central circular feature. The side view shows the profile of the part with a flange. Dimensions are given in millimeters (mm) and inches (in).

Top View Dimensions:

- Overall width: 114 mm (4 1/2 in)
- Overall height: 114 mm (4 1/2 in)
- Inner width: 86 mm (3 3/8 in)
- Inner height: 86 mm (3 3/8 in)
- Central circular feature diameter: 50 mm (2 in)
- Distance from center to corner: 38 mm (1 1/2 in)
- Distance from center to side hole: 38 mm (1 1/2 in)
- Distance from center to bottom hole: 38 mm (1 1/2 in)
- Distance from center to top hole: 38 mm (1 1/2 in)
- Distance from center to side hole (inner): 25 mm (1 in)
- Distance from center to bottom hole (inner): 25 mm (1 in)
- Distance from center to top hole (inner): 25 mm (1 in)
- Distance from center to side hole (outer): 50 mm (2 in)
- Distance from center to bottom hole (outer): 50 mm (2 in)
- Distance from center to top hole (outer): 50 mm (2 in)

Side View Dimensions:

- Overall height: 114 mm (4 1/2 in)
- Overall width: 114 mm (4 1/2 in)
- Inner width: 86 mm (3 3/8 in)
- Inner height: 86 mm (3 3/8 in)
- Central circular feature diameter: 50 mm (2 in)
- Distance from center to corner: 38 mm (1 1/2 in)
- Distance from center to side hole: 38 mm (1 1/2 in)
- Distance from center to bottom hole: 38 mm (1 1/2 in)
- Distance from center to top hole: 38 mm (1 1/2 in)
- Distance from center to side hole (inner): 25 mm (1 in)
- Distance from center to bottom hole (inner): 25 mm (1 in)
- Distance from center to top hole (inner): 25 mm (1 in)
- Distance from center to side hole (outer): 50 mm (2 in)
- Distance from center to bottom hole (outer): 50 mm (2 in)
- Distance from center to top hole (outer): 50 mm (2 in)

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

[illegible]

ADVANCED INFORMATION

January 1990

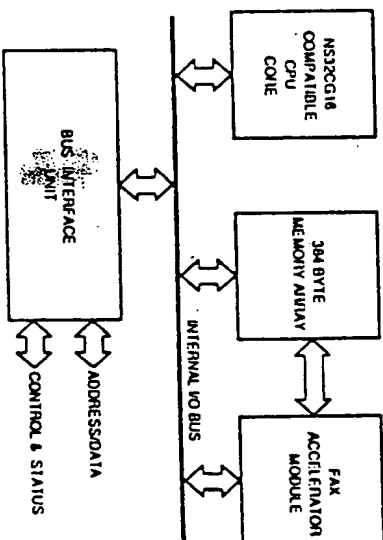
General Description

The NS32FX16 is a high-performance 32-bit Embedded System Processor that is optimized for Group 2 and Group 3 Facsimile applications. Data Modems, Voice Mail systems and Laser Printers. It performs all the computations and control functions required for a stand-alone FAX system, a PC add-in required for a stand-alone FAX system. The FAX/Data modem card or a Laser/FAX system. The NS32FX16 can operate, in real time, V.29 (9600 bps and 7200 bps), V.27 (4800 bps and 2400 bps), and V.21. The NS32FX16 incorporates four main modules: the NS32C6G16 compatible CPU Core, a 384-Byte Memory Array, a FAX Accelerator Module and a Bus Interface Unit.

The CPU Core incorporates a full 32-bit ALU and 32-bit internal data bus. This processor also supports a 16-MbByte linear address space, a 16-bit external data bus and an 8-byte prefetch queue.

The FAX Accelerator Module (FAM) executes vector operations on complex variables and is optimized for modern applications. It is designed to enhance DSP performance on modern Digital Signal Processing (DSP) primitives while preserving the CPU core's structure and programming model. The vector-

Block Diagram



Features

- 32 bit architecture and implementation
- 16-Mbyte linear addressing space
- On-chip FXA Accelerator Module for DSP support
- Special support for graphics applications
 - 18 graphics instructions
 - Efficient fonts & pattern handling
 - Interface to an external BitBLT T processing units for fast color BitBLT operations
- 384-byte on-chip L2M array
- Operating frequency: 15, 20, and 25 MHz
- Binary compatible with the Series 32000 or family
- Floating point support via the NS32081 or NS32081
- Power save mode
- Double-metal CMOS technology
- 68 pin PLCC package
- On-chip clock generator

Embedded System Processor 3, Series 32000 and Tri-STATE are registered trademarks of National Semiconductor Corporation. Patented in is a trademark of Adaptec Systems Inc.

operations can be used to efficiently implement FIR filters and other DSP primitives.

The FAM is attached to the CPU core via the internal I/O Bus. It is treated as a memory mapped I/O device, occupying a reserved memory space. The CPU controls this module via a set of memory mapped registers. The module reduces the load of the main processor by fetching operands using its own address generator. In order to save bus bandwidth, the FAM stores the coefficients of the various filters in an internal 384-byte Memory Array. The 384-byte Memory Array is a shared resource and is usable by both the FAM and the CPU core.

Besides the highly efficient architecture and the addition of the FAM, the NS323FX16 supports all the NS32C616 instructions including graphic enhancements like BitBlit (84-aligned Bit-dot transfer) operations and other special graphic instructions. These graphic enhancements can be used to support Postscript™ applications such as printers and laser/FAX machines.

The microprocessor is packed in a 68-pin Plastic Leaded Chip Carrier (PLCC) package.

Product Introduction

The NS32FX16 is a high speed CMOS microprocessor in National's Embedded System Processor family. It includes two main execution units: the NS32CG16 compatible CPU Core and the FAX Accelerator Module. The CPU Core is designed for general purpose computations and system control functions. The FAX Accelerator Module is tuned to perform the DSP primitives needed in Voice Band Modems. The NS32FX16 also incorporates a 384-byte Memory Array as a shared resource for both the CPU core and the FAX Accelerator Module. Those features make the NS32FX16 an optimal solution for applications in which both general purpose and DSP computations are needed. Such applications include FAX Modems, Voice Compression, and Voice Mail systems.

The NS32FX16 is software-compatible with all other CPUs in the family. The device incorporates all of National's Embedded System Processor advanced architectural features, with the exception of the virtual memory capability.

Brief descriptions of the NS32FX16 features that are shared with other members of the family are provided below:

Powerful Addressing Modes. Nine addressing modes available to all instructions are included to access data structures efficiently.

Data Types. The architecture provides for numerous data types, such as byte, word, doubleword, and BCD, which may be arranged into a wide variety of data structures.

Symmetric Instruction Set. While avoiding special case instructions that compilers can't use, National's Embedded System Processor family incorporates powerful instructions for control operations, such as array indexing and external procedure calls, which save considerable space and time for compiled code.

Memory-to-Memory Operations. National's Embedded System Processor CPUs represent two-address machines. This means that each operand can be referenced by any one of the addressing modes provided.

This powerful memory-to-memory architecture permits memory locations to be treated as registers for all useful operations. This is important for temporary operands as well as for context switching.

Large, Uniform Addressing. The NS32FX16 has 32-bit address pointers that can address up to 16 Mbytes of external memory without any segmentation; this addressing scheme provides flexible memory management without added-on expense.

Modular Software Support. Any software package for National's Embedded System Processor family can be developed independent of all other packages, without regard to individual addressing. In addition, ROM code is totally relocatable and easy to access, which allows a significant reduction in hardware and software costs.

Software Processor Concept. National's Embedded System Processor architecture allows expansions of the instruction set that can be executed by Floating Point slave processors, acting as extensions to the CPU. Current Floating Point slave processors of the Embedded System Processor family are the NS32081 and the NS32381.

To summarize, the architectural features cited above provide three primary performance advantages and characteristics:

- High-Level Language Support
- Easy Future Growth Path
- Application Flexibility

Table of Contents

1.0 PRODUCT INFORMATION	3.4.7 Internal Cycles
1.1 NS32FX16 Special Features	3.4.8 Initiated by On-Chip DMA Controller
2.0 ARCHITECTURAL DESCRIPTION	3.4.9 Slave Processor Communication
2.1 Register Set	3.4.0.1 Slave Processor Bus Cycles
2.1.1 General Purpose Registers	3.4.0.2 Slave Operand Transfer Sequences
2.1.2 Address Registers	3.5 Bus Access Control
2.1.3 Processor Status Register	3.6 Instruction Execution and Status
2.1.4 Configuration Register	3.7 Exception Processing
2.2 Memory Organization	3.7.1 Exception Acknowledge Sequence
2.2.1 Addressing Mapping	3.7.2 Returning from an Exception Service Procedure
2.2.2 Dedicated Tables	3.7.3 Maskable Interrupts
2.3 Instruction Set	3.7.3.1 Non-Vectored Mode
2.3.1 General Instruction Format	3.7.3.2 Vectored Mode: Non-Cascaded Case
2.3.2 Addressing Modes	3.7.3.3 Vectored Mode: Cascaded Case
2.3.3 Instruction Set Summary	3.7.4 Non-maskable Interrupt
2.4 Graphics Support	3.7.5 Traps
2.4.1 Frame Buffer Addressing	3.7.6 Instruction Tracing
2.4.2 BitBLT Fundamentals	3.7.7 Priority Among Exceptions
2.4.2.1 Frame Buffer Architecture	3.7.8 Exception Acknowledge Sequences:
2.4.2.2 Bit Alignment	Detail Flow
2.4.2.3 Block Boundaries and Destination Masks	3.7.8.1 Maskable/Non-Maskable Interrupt Sequence
2.4.2.4 BitBLT Directions	3.7.8.2 Trap Sequence: Traps Other Than Trace
2.4.2.5 BitBLT Variations	3.8 Slave Processor Instructions
2.4.3 Graphics Support Instructions	3.8.1 Slave Processor Protocol
2.4.3.1 BitBLT (Bit-aligned Block Transfer)	3.8.2 Floating-Point Instructions
2.4.3.2 Pattern Fill	
2.4.3.3 Data Compression, Expansion and Magnify	
2.4.3.3.1 Magnifying Compressed Data	
2.5 FAX Accelerator Module	
2.5.1 FAM Operations	
2.5.1.1 Complex Number Representation	
2.5.1.2 Mac Operation	
2.5.1.3 Instruction Set	
2.5.1.4 Circular Buffers	
2.5.1.5 Performance Considerations	
2.5.2 FAM Registers and RAM Array	
2.5.2.1 IAM Coefficient Array C[0]-C[95]	
2.5.2.2 Multiplier Input Register Y	
2.5.2.3 Accumulator A	
2.5.2.4 Data Pointer DP1R	
2.5.2.5 Coefficient Memory Vector Pointer CPTR	
2.5.2.6 Control Register CIL	
2.5.2.7 Status Register ST	
3.0 FUNCTIONAL DESCRIPTION	
3.1 Power and Grounding	
3.2 Clocking	
3.2.1 Power Save Mode	
3.3 Resolving	
3.4 Bus Cycles	
3.4.1 Bus Status	
3.4.2 Basic Read and Write Cycles	
3.4.3 Cycle Extension	
3.4.4 Data Access Sequences	
3.4.4.1 Bit Accesses	
3.4.4.2 Bit Field Accesses	
3.4.4.3 Extending Multiple Accesses	
3.4.5 Instruction Fetches	
3.4.6 Interrupt Control Cycles	
Appendix A: Instruction Formats	
Appendix B: NS32FX16 Instruction Timing	
B.1 Introduction	
B.2 Assumptions	
B.2.1 Definitions	
B.2.2 Interpreting the Table	
B.2.3 Calculating the Effects of Shift Values	
B.2.4 Calculating the Effects of Wait States	
B.3 NS32FX16 General Instruction Timing	
B.3.1 Assumptions	
B.3.2 Definitions	
B.3.3 Calculation of Total Execution Time (TEX)	
B.3.4 Notes on Table Use	
B.3.5 Example of Table Usage	
B.3.6 NS32081 Floating-Point Execution Times	
B.4 FAX Accelerator Module Performance	
B.4.1 Assumptions	
B.4.2 Definitions	
B.4.3 FAM Performance in Clock Cycles	

NS32FX16 Internal CPU Core Registers.....	2-1
Processor Status Register (PSR).....	2-2
Configuration Register (CR).....	2-3
NS32FX16 Memory Organization.....	2-4
Module Descriptor Format.....	2-5
A Sample Link Table.....	2-6
General Instruction Format.....	2-7
Index Byte Format.....	2-8
Displacement Encodings.....	2-9
Correspondence between Linear and Cartesian Addressing.....	2-10
32-Pixel by 32-Scan Line Frame Buffer.....	2-11
Overlapping BiBLT Blocks.....	2-12
BiBLT Instructions Format.....	2-13
BITWT Instruction Format.....	2-14
EXIBLT Instruction Format.....	2-15
MOVWPI Instruction Format.....	2-16
TBITIS Instruction Format.....	2-17
SBITPS Instruction Format.....	2-18
Bus Activity for a Simple BiBLT Operation.....	2-19
FAX Accelerator Module Block Diagram.....	2-20
Memory Organization of a Complex Vector.....	2-21
Power and Ground Connections.....	2-22
Crystal Interconnections - 30 MHz.....	3-1
Crystal Interconnections - 40 MHz, 50 MHz.....	3-2(a)
Recommended Reset Connections.....	3-2(b)
General Reset Timing.....	3-3
Power-on Reset Requirements.....	3-4
Bus Connections.....	3-5
On-Chip Read Cycle Timing.....	3-6
On-Chip Write Cycle Timing.....	3-7
Extension of an On-Chip Read Cycle.....	3-8
Memory Interface.....	3-9
On-Chip Read Cycle.....	3-10
On-Chip Write Cycle.....	3-11(a)
DMAC Initiated Bus Cycle.....	3-11(b)
System Connection Diagram with the NS32081 FPU.....	3-12
System Connection Diagram with the NS32081 FPU.....	3-13(a)
Slave Processor Read Cycle.....	3-13(b)
Slave Processor Write Cycle.....	3-14
HOLD Timing, Bus Initially Idle.....	3-15
HOLD Timing, Bus Initially Not Idle.....	3-16
Interrupt Dispatch and Cascade Tables.....	3-17
Exception Acknowledge Sequence.....	3-18
Return from Trap (RETT) Instruction Flow.....	3-19
Return from Interrupt (RETI) Instruction Flow.....	3-20
Interrupt Control Unit Connections (16 Levels).....	3-21
Cascaded Interrupt Control Unit Connections.....	3-22
Service Sequence.....	3-23
Slave Processor Protocol.....	3-24
Slave Processor Status Word Format.....	3-25
Connection Diagram.....	3-26
Timing Specifications Standard (Signal Valid After Clock Edge).....	4-1
Timing Specifications Standard (Signal Valid Before Clock Edge).....	4-2
Read Cycle.....	4-3
Write Cycle.....	4-4
On-Chip and On-Chip Read Cycles.....	4-5
DMAC Initiated Bus Cycle.....	4-6
	4-7

B.4.4 FAX ACCELERATOR MODULE PERFORMANCE

B.4.1 Assumptions

The FAM instruction execution starts with the CPU core writing to the CTTL register. The execution time is counted from T3 of this transaction until all the results are ready, either in the Accumulator or in the Coefficient array.

It is assumed that there are:

- No wait states in FAM initiated read cycles
- No external HOLD request during the FAM operation for VCMAD, VCMUL and VCMAC instructions.

B.4.2 Definitions

N Number of elements in complex vector
TFAM Number of clock cycles to execute the instruction

B.4.3 FAM Performance in Clock Cycles

FAM instruction	TFAM
VCMAD	9 + (N * 8)
VCMUL	9 + (N * 8)
VCMAC	6 + (N * 8)
VCMAG	5 + (N * 8)

Appendix B: NS32FX16 Instruction Timing (Continued)

8 NS32081 Floating-Point Execution Times
 Following section gives execution timing information for Floating-Point Instructions. Some additional timing
 conditions are used, as given below:

The floating-point operation length:

Standard Floating (32 bits): $T = 1$

Long Floating (64 bits): $T = 2$

The time required to transfer 32 bits of a floating-point value to or from the NS32081 Floating-Point Unit:

$T_1 = 4$ always

The time required to transfer an integer value to or from the NS32081 Floating-Point Unit:

Byte: $T_1 = 2$

Word: $T_1 = 2$

Double-word: $T_1 = 4$

TABLE B-5. Floating-Point Instruction Execution Times

EMONIC	CASE	TEA	TOPD	TOPI	T1	T1	TCY
V1	<AAB>	2	21	-	-	21	23
	<ARB>	-	-	-	-	-	27
	<ARB>	1	1	-	-	1	23
	<ARB>	1	1	-	-	1	27
X, SUBI	<AAB>	2	31	-	-	31	70
	<ARB>	-	-	-	-	-	74
	<ARB>	1	1	-	-	1	70
	<ARB>	1	21	-	-	21	70
J	<AAB>	2	31	-	-	31	30+141
	<ARB>	-	-	-	-	-	30+141
	<ARB>	1	1	-	-	1	30+141
	<ARB>	1	21	-	-	21	30+141
I	<AAB>	2	31	-	-	31	55+301
	<ARB>	-	-	-	-	-	55+301
	<ARB>	1	1	-	-	1	55+301
	<ARB>	1	21	-	-	21	55+301
J, NEGI	<AAB>	1	21	-	-	21	20
	<ARB>	-	-	-	-	-	24
	<ARB>	1	1	-	-	1	20
	<ARB>	1	21	-	-	21	24
Y	<AAB>	1	21	-	-	21	45
	<ARB>	-	-	-	-	-	49
	<ARB>	1	1	-	-	1	45
	<ARB>	1	21	-	-	21	45
LF	<AAB>	1	3	-	-	3	23
	<ARB>	-	-	-	-	-	27
	<ARB>	1	1	-	-	1	23
	<ARB>	1	2	-	-	2	27
/FL	<AAB>	1	3	-	-	3	22
	<ARB>	-	-	-	-	-	26
	<ARB>	1	1	-	-	1	22
	<ARB>	1	2	-	-	2	26
/H	<AAB>	1	1	-	-	1	53
	<ARB>	-	-	-	-	-	53
	<ARB>	1	1	-	-	1	53
	<ARB>	1	1	-	-	1	53
AND, TRUNC, OR	<AAB>	1	1	-	-	1	66
	<ARB>	-	-	-	-	-	13
	<ARB>	1	1	-	-	1	13
	<ARB>	1	1	-	-	1	18

List of Illustrations (Continued)

HOLD Acknowledge (Bus Initially Not Idle).....	4-8
HOLD Acknowledge (Bus Initially Idle).....	4-9
Slave Processor Write Timing.....	4-10
Slave Processor Read Timing.....	4-11
SFC Timing.....	4-12
Relationship of PFS to Clock Cycles.....	4-13
Interlocked Bus Cycle.....	4-14
Clock Waveforms.....	4-15
Power-On Reset.....	4-16
Non-Power-On Reset.....	4-17
NMI Interrupt Signal Timing.....	4-18
INT Interrupt Signal Detection.....	4-19

List of Tables

NS32FX16 Addressing Modes.....	2-1
NS32FX16 Instruction Set Summary.....	2-2
'op' and 'v' Field Encodings.....	2-3
FAX Accelerator Instruction Set Summary.....	2-4
Circular Buffer Size.....	2-5
FAM Register Address Map.....	2-6
External Oscillator Specifications.....	3-1
Bus Cycle Categories.....	3-2
Access Sequences.....	3-3
Interrupt Sequences.....	3-4
Floating-Point Instruction Protocols.....	3-5
Average Execution Time.....	B-1
Shift > 0.....	B-2
Average Execution with Wait States.....	B-3
Instruction Timing Parameters.....	B-4
Floating-Point Instruction Execution Times.....	B-5

1.1 NSJ2FX10 SPECIAL FEATURES

The most relevant of those features are the enhanced Digital Signal Processing performance, which makes the chip very attractive for usage in facsimile, and the graphics support capabilities, that can be used in applications such as printers, CRT terminals, and other varieties of display systems, where text and graphics are to be handled.

The NS22FX16 allows systems to be built with a relatively small amount of random logic. With the external DMA support, the bus can be highly optimized to allow simple interfacing to a large variety of RAMs and peripheral devices. All the relevant bus access signals and clock signals are generated on-chip. The cycle extension logic is also incorporated on-chip.

The device is fabricated in a low-power, double metal, CMOS technology. It also includes a power-save feature that allows the clock to be slowed down under software control, thus minimizing the power consumption. This feature can be used in those applications where power saving during periods of low performance demand is highly desirable.

The bus characteristics and the power save feature are described in the “Functional Description” section. A description of the FPA Accelerator Module is provided in Section 2.5. A general overview of DMI1 operations and a description of the graphics support instructions is provided in Section 2.4. Details on all the NS32C616 Instructions can be found in the NS32C616 Primer/Disassembly Processor Programmer’s Reference Supplement and the related NS32C616 supplement.

Below is a summary of the instructions that are directly applicable to graphics along with their intended use.

Application

INLAND
RAIFOR
RIORI
RIIXON

The **RIIDL T** group of instructions provide a method of quickly imaging characters, creating patterns, windowing and other block oriented effects.

EXHLT
MOVMP Move Multiple Pattern is a very fast

TDITS

Test Bit String will measure the length of 1's or 0's in an image, supporting many data compression methods (RL). TDITS may also be used to test for boundaries of images.

Application

SBIT5

Set Bit String is a very fast instruction for filling objects, outline characters and drawing horizontal lines. The SBIT5 and SBIT5 instructions support the CCIT standard for compression algorithms used by Group 3 and Group 4 facsimile machines.

SDBITS

Set Bit Perpendicular String is a very fast instruction for drawing vertical, horizontal, and 45° lines. In printing applications SDBITS and SDBTIPS may be used to express portrait and landscape respectively from the same compressed font data. The size of the character may be scaled as it is drawn.

SBIT
CBIT
TBIT
MBIT

The Bit Group of instructions enable single pixels anywhere in memory to be set, cleared, tested or inverted.

1511

INDEX

The **INDEX** instruction combines a multiply-add sequence into a single instruction. This provides a fast translation of an X-Y address to a pixel relative address.

2.0 Architectural Description

2.1 REGISTER SET

The NS32FX16 CPU core has 17 internal registers grouped according to function as follows: 8 general purpose, 7 address, 1 processor status and 1 configuration. *Figure 2-1* shows the NS32FX16 internal registers of the CPU core.

Besides the CPU core registers, the NS32FX16 also has 6 registers, FAX Accelerators and 384 byte TMM which can be accessed as memory-mapped I/O. Refer to section 2.5 for more details.

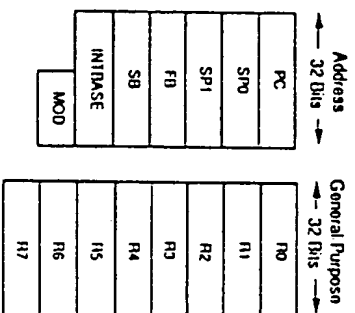


FIGURE 2-1. NS32FX16 Internal CPU Core Registers

TABLE B-4. Instruction Timing Parameters

MNEMONIC	TEA	TOPB	TOPW	TOPO	TOPI	TCV	L	NOTES
INSTR	2	-	-	2	1	29-39	-	load in memory
INSTR	1	-	-	-	1	28-96	-	load in register
INSTR	2	-	-	2	1	38-49	-	
JUMP	1	-	-	1	1	5-15	-	
JUMP	1	-	-	-	-	2-6	-	
LEIGH	1	-	-	-	1	19-33	-	
LEIGH	2	1	-	-	2	14-45	-	
LEIGH	2	-	-	-	4	23	16	
LEIGH	2	-	-	-	3	54-73	16	<M><M1><R1>
MOVW	2	-	-	-	2	1-3	-	n = # of elements in block
MOVW	2	-	-	-	2	3-20	-	<M><M1><R1>
MOVW	1/0	-	-	-	1/0	2/3	-	<M><R1>
MOVW	-	-	-	-	-	-	-	See graphics instructions
MOVW	-	-	-	-	2	24-54	-	B, W and/or M option in effect
MOVW	-	-	-	-	-	-	-	Translated
MOVW	-	n	-	-	2	27-54	-	
MOVW	2	1	-	1	-	6	-	
MOVW	2	1	1	-	-	6	-	
MOVW	2	-	1	1	-	6	-	
MOVW	2	1	-	1	-	5	-	
MOVW	2	1	1	-	-	5	-	
MOVW	2	-	1	-	-	5	-	
MOVW	2	-	1	1	-	5	-	
MUL	2	-	-	-	3	15	16	
NEG	2	-	-	-	2	5	16	
NEG	2	-	-	-	-	3	-	
NOT	2	-	-	-	2	5	-	
OR	2	-	-	-	3	3-4	-	<M><M1><R1>
OR	2	-	-	-	3	49-55	16	
OR	2	-	-	-	3	57-62	16	
RESTORE	-	-	-	n	-	5-12	-	n = # of general registers restored
RESTORE	-	-	-	-	-	2-8	-	
RESTORE	-	1	3	3	-	39-45	-	
RESTORE	-	2	2	-	-	35-41	-	
RESTORE	2	1	-	-	2	14-45	-	
RESTORE	-	-	1	2	-	2-6	-	
SAVE	1	-	-	n	-	4-13	-	False/True saved
SAVE	-	-	-	-	-	-	-	n = # of general registers saved
SAVE	2	2	-	-	1	15-7	-	<M><R1>
SAVE	2	2	-	1	15-7	-	-	<M><R1>
SAVE	-	-	-	-	-	15	-	
SAVE	-	-	-	n	-	27-51	-	n = # of elements, not translated
SAVE	-	n	-	n	-	30-51	-	Translated
SAVE	1	-	-	-	1	21-27	-	<M><M1><R1>
SAVE	2	-	-	-	3	3-4	-	<M><R1>
SAVE	2	-	-	-	3	16-18	-	no carry
SAVE	-	-	4	3	-	40	-	<M><R1>
SAVE	2	-	-	-	1	1-4	-	7 - until an interrupt/reset
WAIT	-	-	-	-	-	6-7	-	<M><M1><R1>
WAIT	2	-	-	-	3	4-4	-	<M><M1><R1>

TABLE B-4. Instruction Timing Parameters

INSTRONIC	TEA	TOPB	TOPW	TOPD	TOPI	TCY	L	NOTES
INSI	2	-	-	-	2	9/8	-	3IC <0>3IC = 0
CIU	1	-	-	-	2	16/15%21	-	<MI>, no branch/branch
CIU	-	-	-	-	-	18/17%32	-	, no branch/branch
DIJ	21/10	-	-	-	31/10	3/4/4	-	<MI> <MI> <RI>
DIJ	21/10	-	-	-	31/10	3/4/4	-	<MI> <MI>
DIJ	2	-	-	-	3	16/18	-	no carry/carry
DIJ	1/0	-	-	-	2/0	6/4	-	<MI>
DIJ	2/1	-	-	1/0	-	2/3	-	<MI>
DIJ	1	-	-	-	1	6	-	<MI> <MI>
DIJ	21/10	-	-	-	31/10	3/4/4	-	<MI> <MI> <RI>
SH	2	1	-	-	2	14/45	-	
ISD	-	-	-	-	-	7/6%10	-	no branch/branch
ISD	21/10	-	-	-	31/10	3/4/4	-	<MI> <MI>
ISD	1	1	-	-	-	18%22	-	
ISD	1	1	-	-	-	30%34	-	
ISD	1	1	-	-	-	18%22	-	
ISD	1	-	1	-	-	30%34	-	
PT	-	-	4	3	-	40	-	
SI	1	-	-	-	1	4%0	-	
SI	-	-	-	-	-	6%16	-	
ASEI	1	-	-	-	1	4%9	-	<MI> <RI>
DIJ	2/1	2/0	-	-	1	15/7	-	<MI>
DIJ	2/1	2/0	-	-	1	15/7	-	high/low/low
HEC	2	-	-	-	3	71/0/11	-	<MI> <MI> <RI>
HEC	2	-	-	-	21/10	3	-	n = # of elements in block
MI	21/10	-	-	-	21/10	91/124	-	n = # of elements in block
MI	2	-	-	-	21/10	3	-	<MI>
MI	1/0	-	-	-	1/0	3	-	n = # of elements, not translated
MI	-	-	-	-	21/10	351/153	-	n = # of elements, not translated
MI	-	-	-	-	21/10	381/153	-	translated
MI	2	-	-	-	2	7	-	
MI	2	-	-	-	1	7	-	
MI	-	-	-	-	-	16%21	-	
MI	-	-	-	-	3	13%18	-	
MI	1	-	-	-	3	5/1	-	<MI>
MI	2/1	-	-	-	-	3%7	-	
MI	-	-	-	-	-	58/68	16	n = # of general registers
MI	2	-	-	-	3	41/118	-	n = # of general registers
MI	-	-	-	-	n+1	-	-	saved
MI	-	-	-	-	n+1	-	-	n = # of general registers
MI	-	-	-	-	n+1	-	-	restored
MI	-	-	-	-	n+1	-	-	n = # of general registers
MI	2	-	-	-	1	19/29	-	held in memory
MI	2	-	-	-	1	17/51	-	held in register
MI	2	-	-	-	1	26/36	-	
MI	2	2	-	-	1	24/28	24	
MI	-	-	0/4	0/3	-	6/44	-	no trap/carry
MI	2/1	2/0	-	-	1	17/9	-	<MI>
MI	2	-	-	-	2	25	16	

2.1.1.1 General Purpose Registers

2.1.1.2 Address Registers

PC—Program Counter. The PC register is a pointer to the first byte of the instruction currently being executed. The PC is used to reference memory in the program section.

When a reference is made to the selected Stack Pointer (see PSR S bit), the terms 'SP Register' or 'SP' are used. SP refers to either SP0 or SP1, depending on the setting of the S bit in the PSR register. If the S bit in the PSR is 0, SP refers to SP0. If the S bit in the PSR is 1 then SP refers to SP1.

FP—Frame Pointer. The FP register is used by a procedure to access parameters and local variables on the stack. The FP register is set upon procedure entry with the ENTER instruction and restored on procedure termination with the EXIT instruction.

SB—Static Base. The SB register points to the global variables of a software module. This register is used to support relocatable global variables for software modules. The SB register holds the lowest address in memory occupied by the global variables of a module.

INTBASE--Interrupt Base. The INTBASE register holds the address of the dispatch table for interrupts and traps (Section 3.7.1).

MOD—Module. The MOD register holds the address of the module descriptor of the currently executing software module. The MOD register is 16 bits long; therefore, the module table must be contained within the first 64 Kbytes of memory.

2.1.1.3 Processor Status Register

executing in Supervisor Mode.

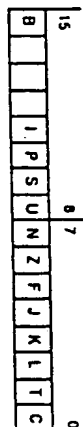


FIGURE 2-2. Processor Status Register (PSR)

selling at 0 (in case of common) or 1 (very rare borrow).

to 1. A MAC trap is executed when every instruction (Section 3.7.6):

at loading-point components, and are not cleared.

The F bit is a general condition flag, which is reserved for use by the CIO.

arithmetic instructions use it to indicate overflow).

1- if the second operand is equal to the first

N The N bit is altered by comparison instructions. In a comparison instruction the N bit is set to

operand, when both operands are interpreted as signed integers. Otherwise, it is set to "0".

be executed. If the U bit is "0" then all instructions may be executed. When U = 0 the

intermediate, "flag" bit. The **Supervisor Mode** processor is said to be in **Supervisor Mode** when $U = 1$. The processor is said to be in **User Mode**. A **User Mode** program is restricted from executing certain instructions and from accessing certain registers which could interfere with the operating system. For example, a **User Mode** program is prevented from changing the setting of the flag used to indicate its own privilege mode.

2.2 MEMORY ORGANIZATION

- A Supervisor Mode program is assumed to be a trusted part of the operating system, hence it has no such restrictions.
- The S bit specifies whether the SPO register or SPT register is used as the Stack Pointer. The bit is automatically cleared on interrupts and traps. It may have a setting of 0 (use the SPO register) or 1 (use the SPT register).
- The P bit prevents a TRC trap from occurring more than once for an instruction (Section 3.7.6). It may have a setting of 0 (no trace pending) or 1 (trace pending).
- If I=1, then all interrupts will be accepted. If I=0, only the NMI interrupt is accepted. Trap enables are not affected by this bit.
- Reserved for use by the CPU. This bit is set to 1 during the execution of the EXTLT instruction and causes the BPU signal to become active. Upon reset, B is set to zero and the BPU signal is set high.

Note 1: When an interrupt is acknowledged, the 0, 1, P, S and U bits are set to zero and the BPU signal is set high. A return from interrupt will restore the original values from the copy of the PSR register saved in the interrupt stack.

Note 2: If BOUT (B8) instructions are executed in an interrupt routine, the PSH bit 3 and it must be cleared first.

2.1.4 Configuration Register

The Configuration Register (CFG) is 8 bits wide, of which 4 bits are implemented. The implemented bits enable various operating modes for the CPU, including execution of floating-point instructions, processing of exceptions and selection of clock scaling factor. CFG is programmed by the SEICFG instruction. The format of CFG is shown in Figure 2-3. The various control bits are described below.

7							0
RES	C	M	F	I			

FIGURE 2-3. Configuration Register (CFG)

- I Interrupt vectoring. This bit controls whether maskable interrupts are handled in nonvectored (I=0) or vectored (I=1) mode. Refer to Section 3.7.3 for more information.
- F Floating-point instruction set. This bit indicates whether a floating-point unit (FPU) is present to execute floating-point instructions. If this bit is 0 when the CPU executes a floating-point instruction, a Trap (UND) occurs. If this bit is 1, then the CPU transfers the instruction and any necessary operands to the FPU using the slave-processor protocol described in Section 3.8.1.
- M Clock scaling. This bit is used in conjunction with the C bit to select the clock scaling factor.
- C Clock scaling. Same as the M bit above. Refer to Section 3.2.1 on "Power Save Mode" for details.

7							0
							A

Byte at Address A

Two contiguous bytes are called a word. Except where noted, the least significant byte of a word is stored at the lower address, and the most significant byte of the word is stored at the next higher address. In memory, the address of a word is the address of its least significant byte, and a word may start at any address.

15							0
							A+1
							A

Word at Address

Two contiguous words are called a double word. Except where noted, the least significant word of a double word is stored at the lowest address and the most significant word of the double word is stored at the address two higher. In memory, the address of a double word is the address of its least significant byte, and a double word may start at any address.

31							0
							A+3
							A+2
							A+1
							A

Double Word at Address A

Although memory is addressed as bytes, it is actually organized as words. Therefore, words and double words that are aligned to start at even addresses (multiples of two) are accessed more quickly than words and double words that are not so aligned.

B.3.3 Calculation of Total Execution Time (TEX)

- TEX is obtained by performing the following steps:
- Find the desired instruction in the table.
 - Calculating the values of TEA, TOPI, etc., using the numbers in the table and the equations given on the preceding page.
 - The result derived by adding together these values is the execution time (TEX) in clock cycles.

B.3.4 Notes on Table Use

Values in the TEA column (see Tables B-4 and B-5) indicate the number of effective addresses to be calculated. If the value in this column is less than the number of general operands in the instruction, this is because one or both operands are in registers and the instruction has an optimized form which eliminates TEA for such operands.

In the L column, multiply the entry by the operation length in bytes (1, 2 or 4).

In the TCY column, special notations sometimes appear:

n1-n2 means n1 minimum, n2 maximum.
 n1%2 means that the instruction flushes the instruction queue after n1 clock cycles and nonsequentially fetches the next instruction. The value n2, indicating the total number of clock cycles in internally executing the instruction (including n1), is not generally useful. The most accurate technique for determining such timing depends on the size and alignment of the basic instruction portion of the next instruction, plus index bytes. If this portion can be read in one memory cycle, then the execution time is n1+10 (including the memory cycle). If more memory cycles are required, the value is n1+5+4*m, where m is the number of memory cycles required.

In the Notes column, notations held within angle brackets <> indicate alternatives in the form of the instruction which affect the execution time. A table entry which is affected by the form of the instruction may have multiple values, separated by slashes, corresponding to the alternatives. The notations are:

<M>	Memory form
<R>	Register form
<M>	Memory-to-Memory form
<M>	Memory-to-Register form
<R>	Register-to-Register form
<R>	Register-to-Memory

B.3.5 Example of Table Usage

Calculate TEX for the instruction: CMPW R0, TOS.

Operand A is in a register. Operand B is in memory. The table values must be used corresponding to the <M> case as given in the Notes column (<M> meaning "anything to memory").

Only the TEA, TOPI and TCY columns have values assigned for the CMPW instruction; therefore, they are the only ones that need to be calculated to find TEX. The blank columns are irrelevant to this instruction.

The TEA column contains 2 for the <M> case. This means that effective address times have to be calculated for both operands. (For the <M> case, the Register operand requires no TEA time; therefore, only the Memory operand TEA is necessary.) From the equations:

$$\begin{aligned} \text{TEA (Register mode)} &= 2 \\ \text{TEA (Top-of-Stack mode, read access data)} &= 2 \\ \text{Total TEA} &= 2 + 2 = 4 \end{aligned}$$

The TOPI column represents potential operand transfers to or from memory. For a Compare instruction, each operand is read once, for a total of two operand transfers

$$\begin{aligned} \text{TOPI (Word, Register)} &= 0 \\ \text{TOPI (Word, TOS)} &= \text{TOPW} = 3 \text{ (assuming aligned, no wait)} \\ \text{Total TOPI} &= 3 \end{aligned}$$

TCY is the time required for internal operation within the CPU. The TCY value for this case is 3.

$$\text{TEX} = \text{TEA} + \text{TOPI} + \text{TCY} = 4 + 3 + 3 = 10 \text{ machine cycles.}$$

If the CPU is running at 15 MHz, then a machine cycle (clock cycle) is 66 ns. Therefore, this instruction takes 10 x 66 ns, or 660 ns to execute.

Appendix B: NS32FX16 Instruction Timing (Continued)

B.3.2 Definitions

TEA The time required to calculate an operand's effective address. For a Register or Immediate operand, this includes the fetch of that operand.

TOPB The time needed to read or write a memory byte.

TOPW The time needed to read or write a memory word.

TOPD The time needed to read or write a memory double-word.

TOPH The time needed to read or write a memory operand, where the operand size is given by the operation length of the instruction. It is always equivalent to either TOPB, TOPW or TOPD.

TCY Internal processing overhead, in clock cycles.

L Internal processing whose duration depends on the operation length. The number of clock cycles is derived by multiplying this value by the number of bytes in the operation length.

Equations

TOPB If operand is in a register or is immediate then TOPB = 0

else TOPB = 3

If operand is in a register or is immediate then TOPW = 0

else if word-aligned (even address) then TOPW = 3

else TOPW = 7

If operand is in a register or is immediate then TOPD = 0

else if word-aligned (even address) then TOPD = 7

else TOPD = 11

If operand is in a register or is immediate then TOPH = 0

else if L = byte then TOPH = TOPB

else if L = word then TOPH = TOPW

else if L = double-word then TOPH = TOPD

TCY L

If (operation length) = byte then L = 1

else if L = word then L = 2

else if L = double-word then L = 4

If REGISTER addressing then TEA = 2

If IMMEDIATE or ABSOLUTE addressing then TEA = 4

If REGISTER RELATIVE or MEMORY SPACE addressing then TEA = 5

If MEMORY RELATIVE addressing then TEA = 7 + TOPD

If TOP OF STACK addressing then

If access class = write then TEA = 4

If access class = read then TEA = 2

else TEA = 3

If EXTERNAL addressing then TEA = 11 + 2 * TOPD

If SCALED INDEXED addressing then TEA = 11 + 112

where 111 depends on scale factor:

If byte indexing 111 = 5

If word indexing 111 = 7

If double-word indexing 111 = 8

If quad-word indexing 111 = 10

and 112 = TEA of the base mode except:

if base mode is REGISTER then 112 = 5

if base mode is TOP OF STACK then 112 = 4

2.0 Architectural Description (Continued)

2.2.1 Addressing Mapping

Besides addressing the 16 Mbyte (24-bit address) external memory space, the NS32FX16 can also address 4 Gbytes (32-bit address) of its on-chip memory space (see Figure 2-4). However, the upper 8 address bits are not issued to the memory or I/O outside the microprocessor. They are used only on-chip to access memory-mapped I/O (I/O and registers). This I/O is mapped in the FFFFD000-FFFFFFF (hex) address range, which is part of the dedicated address space defined for National's Embedded System Processor architecture. The address space FFFFD000-FFFFFFF (hex) is dedicated for the FAX Accelerator.

When accessing external memory, the address bits 24 to 31 (available on-chip only) should be kept zero.

ADDRESS (hex)

FFFFFFF	RESERVED BY NSC
FFFD410	ON-CHIP MEMORY, MAPPED TO (FAX ACCELERATOR AND INTERNAL RAM)
FFFD000	RESERVED BY NSC
01000000	INTERRUPT CONTROL
00FFFFFF	ADDRESSES AVAILABLE OFF-CHIP (MULTIMAX AND I/O)
00FFFC00	
00FFFDFF	
0000000	

FIGURE 2-4. NS32FX16 Memory Organization

2.2.2 Dedicated Tables

Two of the NS32FX16 dedicated registers (MOD and INTBASE) serve as pointers to dedicated tables memory.

The INTBASE register points to the Interrupt Dispatch and Cascade Tables. These are described in Section 3.7.

The MOD register contains a pointer into the Module Table, whose entries are called Module Descriptors. A Module Descriptor contains four pointers, three of which are used by the NS32FX16. The MOD register contains the address of the Module Descriptor for the currently running module. It is automatically updated by the Call External Procedure instructions (CXP and CXPD).

The format of a Module Descriptor is shown in Figure 2-5. The Static Base entry contains the address of static data assigned to the running module. It is loaded into the CPU Static Base register by the CXP and CXPD instructions. The Program Base entry contains the address of the first byte of instruction code in the module. Since a module may have multiple entry points, the Program Base pointer serves only as a reference to find them.

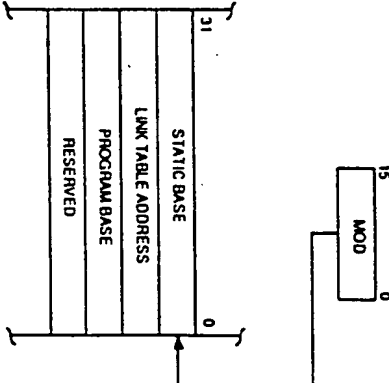


FIGURE 2-5. Module Descriptor Format

The Link Table Address points to the Link Table for the currently running module. The Link Table provides the information needed for:

- 1) Sharing variables between modules. Such variables are accessed through the Link Table via the External addressing mode.
- 2) Transferring control from one module to another. This is done via the Call External Procedure (CXP) instruction.

The format of a Link Table is given in Figure 2-6. A Link Table Entry for an external variable contains the 32-bit address of the variable. An entry for an external procedure contains two 16-bit fields: Module and Offset. The Module field contains the new MOD register contents for the module being entered. The Offset field is an unsigned number giving the position of the entry point relative to the new module's Program Base pointer.

For further details of the functions of those tables, see the Series 32000 Instruction Set Reference Manual.

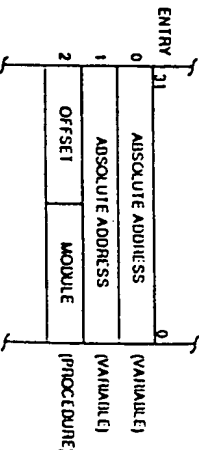


FIGURE 2-6. A Sample Link Table

NUMBER OF CYCLES

Each Displacement field may contain one of two displacements, or one immediate value. The size of a Displacement field is encoded within the top bits of that field, as shown in *Figure 2-9*, with the remaining bits interpreted as a signed (two's complement) value. The size of an immediate value is determined from the size of the field. Both Displacement and Immediate fields are stored most significant byte first.

Note that this is different from the memory representation of data (Section 2.2).

Some instructions require additional "implied" immediates and/or displacements, apart from those associated with addressing modes. Any such extensions appear at the end of the instruction, in the order that they appear within the list of operands in the instruction definition (Section 2.3.3).

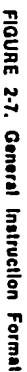


FIGURE 2-7. General Instruction Format

B.3 NS32FX16 GENERAL INSTRUCTION TIMING

B.3.1 Assumptions

The entire instruction, with all displacements and immediate operands, is assumed to be present in the instruction queue when needed.

Interference from instruction prefetches, which is very dependent upon the preceding instruction(s), is ignored. This assumption tends to affect the timing estimate in an optimistic direction.

It is assumed that all memory operand transfers are completed before the next instruction begins execution. In the case of an operand of access

class rmm in memory, this is pessimistic, as the While transfer occurs in parallel with the execution of the next instruction.

It is assumed that there is no overlap between the fetch of an operand and the following sequences of microcode. This is pessimistic, as the fetch of Operand A generally occurs in parallel with the effective address calculation of Operand B, and the fetch of Operand B occurs in parallel with the execution phase of the instruction.

Where possible, the values of operands are taken into consideration when they affect instruction timing, as a range of times is given. Where this is not done, the worst case is assumed.

Appendix B: NS32FX16 Instruction Timing (Continued)

B.2.3 Calculating the Effects of Shift Values

The formulae in the table give the average execution time for the BITLT instructions (BBOR, BBXOR, BBAND, BBFOR, BITWT) with a shift of 0 to 8 bits. The BITWT and BBFOR instructions, however, optimize a shift of 0 bits by reading and writing only a word of the destination data (16 bits). For shifts greater than 0, they still must read and write a double-word of data (32 bits). Note that the EXIBLT instruction is not affected by the shift amount.

Shifts of greater than 8 bits add 1 clock per bit of shift. For a word of data read, for example, the BBOR shift of 15 bits yields the following formula:

$$42 + (107 + 44 \cdot (\text{width} - 2)) \cdot \text{height} + ((\text{shift} - 8) \cdot \text{width} \cdot \text{height})$$

Inserting the previous example of the 10 by 50 BBOR:

$$45 + (107 + 44 \cdot (10 - 2)) \cdot 50 + ((15 - 8) \cdot 10 \cdot 50) \\ \text{or} \\ 42 + (107 + 352) \cdot 50 + (7 \cdot 500) = 26,482 \text{ clocks or} \\ 1.77 \text{ msec @ } 15 \text{ MHz}$$

This represents the "worst case" time for this instruction, since a shift of greater than 15 bits can be handled by moving the source and destination pointers by 2 bytes and adjusting the shift amount.

The "best case" and "average case" times for most instructions are the same, due to reading the destination data during the shifting of the source data. This parallel operation is a feature of National's Embedded System Processor.

The "best case" for the BITWT and BBFOR instructions is a shift of zero bits. This is due to further internal optimization of these instructions, realizing that only a word of the destination data needs to be operated on if a shift of zero is specified.

Table B-2 shows the expected timing information for shifts greater than 8 bits.

TABLE B-2. Shift > 8

INSTRUCTION	NUMBER OF CYCLES
BBOR	$42 + (107 + 44 \cdot (\text{width} - 2)) \cdot \text{height} + ((\text{shift} - 8) \cdot \text{width} \cdot \text{height})$
BBXOR	$44 + (107 + 44 \cdot (\text{width} - 2)) \cdot \text{height} + ((\text{shift} - 8) \cdot \text{width} \cdot \text{height})$
BBAND	$45 + (111 + 44 \cdot (\text{width} - 2)) \cdot \text{height} + ((\text{shift} - 8) \cdot \text{width} \cdot \text{height})$
BBFOR	$48 + (74 + 32 \cdot (\text{width} - 2)) \cdot \text{height} + ((\text{shift} - 8) \cdot \text{width} \cdot \text{height})$
BBSTO	$56 + (170 + 60 \cdot (\text{width} - 2)) \cdot \text{height} + ((\text{shift} - 8) \cdot \text{width} \cdot \text{height})$
BITWT	$28 + (\text{shift} - 8)$
EXIBLT	$35 + (17 + 13 \cdot \text{width}) \cdot \text{height (pre-read)}$
	$35 + (11 + 13 \cdot \text{width}) \cdot \text{height (no pre-read)}$

B.2.4 Calculating the Effects of Wait States

Since the new NS32FX16 instructions make use of the pipelined reads and writes of National's Embedded System Processor, calculation of the effect of wait states is rather difficult. As an example, in the MOVSI instruction group, each wait state on read operations adds 1 clock cycle per read bus access. Each wait state on write operations subtracts 1 clock cycle per write bus access from the ICY of the instruction, since updating the pointers occurs in parallel with the write operation. This means that wait states can be added to write cycles without changing the execution time of the instruction, up to a maximum of 13 wait states on writes for MOVSB and MOVSW, and 4 wait states on writes for MOVSD. At zero wait states, a MOVSD of 1,056,000 bytes (264,000 double-words) executes in:

$$30 + (10 \cdot 264,000) + (2 \cdot 264,000 \cdot 8) \\ \text{or} \\ 30 + 2,640,000 + 4,224,000 = 6,864,030 \text{ clocks or } 458 \\ \text{msec @ } 15 \text{ MHz}$$

With two wait states on read and write (Twaird = 2 and Twairw = 2) — see Table B-3 — the time becomes:

$$30 + (10 \cdot 264,000) + (2 \cdot 264,000 \cdot 8) + 0 + (2 \cdot 2) \cdot 264,000 \\ \text{or} \\ 30 + 2,640,000 + 4,224,000 + 1,056,000 = 7,920,030 \\ \text{clocks or } 528 \text{ msec @ } 15 \text{ MHz}$$

Instructions that have shift amounts, such as BBOR, BBXOR, BBAND, BBFOR and BITWT, make use of the parallel nature of National's Embedded System Processor by doing the actual shift during the reading of the double-word destination data. This means that if there are wait states on read operations, those instructions are able to shift bytes, without impacting the overall time. For example, the total execution time for a BBFOR operation, shifting 8 bits with 2 wait states on read operations, is the same execution time as for a BBFOR operation shifting by 12 bits. This is because a destination word takes 4 clock cycles longer than a no-wait-state double-word read does. Note that this effect is not valid for more than 4 wait states, because at 4 wait states all possible shift values (0 - 15) are "hidden" during the destination read.

Note that in Table B-3, Twairds refers to the reading of the source data, or of the table data used for a particular instruction. Twairw refers to the reading of the destination data, or of the data on which to be operated. Table B-3 shows the expected timing of instructions with wait states. Again, this is the average execution time, using a shift of eight (whole appropriate).

Twairw is equal to (Twairds + 2 + Twairdd + 2 + Twairwrd). This represents one BitLT transfer, which makes the following equations easier.

2.0 Architectural Description (Continued)

7	3	2	0
GEN. ADDR. MODE			REG. NO.

FIGURE 2-8. Index Byte Format

7	
0	SIGNED DISPLACEMENT

Byte Displacement: Range -64 to +63

7	
0	SIGNED WORD DISPLACEMENT

Word Displacement: Range -8192 to +8191

7	
0	SIGNED DOUBLEWORD DISPLACEMENT

Double Word Displacement:
Range (Effective Addressing Space)

FIGURE 2-9. Displacement Encodings

2.2.2 Addressing Modes

The NS32FX16 CPU generally accesses an operand by calculating its Effective Address based on information available when the operand is to be accessed. The method to be used in performing this calculation is specified by the programmer as an "addressing mode." Addressing modes in the NS32FX16 are designed to optimally support high-level language accesses to variables. In nearly all cases, a variable access requires

only one addressing mode, within the instruction that acts upon that variable. Extraneous data movement is therefore minimized.

NS32FX16 Addressing Modes fall into nine basic types:

Register: The operand is available in one of the eight General Purpose Registers. In certain Slave Processor instructions, an auxiliary set of eight registers may be referenced instead.

Register Relative: A General Purpose Register contains an address to which is added a displacement value from the instruction, yielding the Effective Address of the operand in memory.

Memory Space: Identical to Register Relative above, except that the register used is one of the dedicated registers PC, SP, SB or FP. These registers point to data areas generally needed by high-level languages.

Memory Relative: A pointer variable is found within the memory space pointed to by the SP, SB or FP register. A displacement is added to that pointer to generate the Effective Address of the operand.

Immediate: The operand is encoded within the instruction. This addressing mode is not allowed if the operand is to be written.

Absolute: The address of the operand is specified by a displacement field in the instruction.

External: A pointer value is read from a specified entry of the current Link Table. To this pointer value is added a displacement, yielding the Effective Address of the operand.

Top of Stack: The currently-selected Stack Pointer (SP) or SP1 specifies the location of the operand. The operand is pushed or popped, depending on whether it is written or read.

Scaled Index: Although encoded as an addressing mode, Scaled Indexing is an option on any addressing mode except Immediate or another Scaled Index. It has the effect of calculating an Effective Address, then multiplying any General Purpose Register by 1, 2, 4 or 8 and adding into the total, yielding the final Effective Address of the operand.

Table 2-1 is a brief summary of the addressing modes. For a complete description of their actions, see the Series 32000 Instruction Set Reference Manual.

In addition to the general modes, Register-Indirect with auto-increment/decrement and wraps or pitch are available on several of the graphics instructions.

2.0 Architectural Description (Continued)

TABLE 2-1. NS32FX16 Addressing Modes

ENCODING	MODE	ASSEMBLER SYNTAX	EFFECTIVE ADDRESS
Register			
00000	Register 0	R0 or R0	None: Operand is in the specified register.
00001	Register 1	R1 or R1	
00010	Register 2	R2 or R2	
00011	Register 3	R3 or R3	
00100	Register 4	R4 or R4	
00101	Register 5	R5 or R5	
00110	Register 6	R6 or R6	
00111	Register 7	R7 or R7	
Register Relative			
01000	Register 0 relative	disp(R0)	Disp + Register.
01001	Register 1 relative	disp(R1)	
01010	Register 2 relative	disp(R2)	
01011	Register 3 relative	disp(R3)	
01100	Register 4 relative	disp(R4)	
01101	Register 5 relative	disp(R5)	
01110	Register 6 relative	disp(R6)	
01111	Register 7 relative	disp(R7)	
Memory Relative			
10000	Frame memory relative	disp2(disp1 (FP))	Disp2 + Pointer: Pointer found at address Disp 1 + Register. -SP is either SP0 or SP1, as selected in PSR.
10001	Stack memory relative	disp2(disp1 (SP))	
10010	Static memory relative	disp2(disp1 (SB))	
Reserved			
10011	(Reserved for Future Use)		
Immediate			
10100	Immediate	value	None: Operand is input from instruction queue.
Absolute			
10101	Absolute	@disp	Disp.
External			
10110	External	EXT (disp) + disp2	Disp2 + Pointer: Pointer is found at Link Table Entry number Disp1.
Top Of Stack			
10111	Top of stack	TOS	Top of current stack, using either User or Interrupt Stack. Pointer, as selected in PSR. Automatic Push/Pop included.
Memory Space			
11000	Frame memory	disp(FP)	Disp + Register: -SP is either SP0 or SP1, as selected in PSR.
11001	Stack memory	disp(SP)	
11010	Static memory	disp(SB)	
11011	Program memory	* + disp	
Scaled Index			
11100	Index, bytes	mode[Rn:B]	EA (mode) + Rn. EA (mode) + 2 x Rn. EA (mode) + 4 x Rn. EA (mode) + 8 x Rn. "Mode" and "n" are contained within the Index Byte. EA (mode) denotes the effective address generated using mode.
11101	Index, words	mode[Rn:W]	
11110	Index, double words	mode[Rn:D]	
11111	Index, quad words	mode[Rn:Q]	

APPENDIX A. REGISTER AND INSTRUCTION

A.1 INTRODUCTION

This chapter shows the expected timing of the NS32FX16 graphics instructions. Refer to Section B.3.

As this is advance information, it is subject to change without notice.

A.2 ASSUMPTIONS

The cycle time is one T-state of the Series 32000. Equivalent to one-half of the input clock frequency present on the OSCIN pin of the NS32FX16. A 30-MHz clock source, therefore, yields a cycle time of 66.67 ns. Since the C and M bits of the NS32FX16's configuration register control an on-chip clock divider, setting those bits divides the clock by 1, 2, 4 or 8 to give a cycle time (with a 30-MHz clock source) of 66.67, 133.3 ns, 266.7 ns and 533.3 ns. This first section describes timing of the graphics instructions. When needed, the entire instruction is assumed to be present in the instruction queue. Interference from instruction prefetches is ignored, with the exception of the BITWT instruction, where a no-wait-state prefetch is included (four clock cycles).

It is assumed that all memory operand transfers are completed before the next instruction begins execution. In the case of an operand of access class RMW in memory, this is pessimistic, as the write transfer occurs in parallel with the execution of the next instruction.

Where possible, the values of operands are taken into consideration when they affect instruction timing, and a range of times is given. Where this is not done, the average case is assumed. All memory accesses are assumed to be word aligned. Non-word-aligned data is acceptable but causes the execution time of a given instruction to increase.

The variety of definitions that follows allows accurate prediction of system performance when, for example, the source data may be in ROM and the destination may be in RAM, each having a different number of wait states. The number of wait states refers to the number of additional clock cycles requested via the CWAIT or WAIT pins of the NS32FX16, on a given byte or word-memory access.

A.2.1 Definitions

Twailrd
The number of wait states applied for a Read operation.

Twailwr
The number of wait states applied for a Write operation.

Twailrds
The number of wait states applied for a Read operation on source data. This also refers to the number of wait states applied for a table memory access (in the SBITS instruction, for example).

Twailrd
The number of wait states applied for a Read operation on destination data.

Twailwr
The number of wait states applied for a Write operation on destination data.

Twailbll
Twailrds + Twailrd + 2 * Twailwr
The value used for BITLT timing.

width
The width of a BITLT operation, in words.

height
The height of a BITLT operation, in bit lines.

shift
The number of bits of shift applied.
The Average Execution Time table (see Table B-1) is the execution time for the NS32FX16 instructions based on an average shift of eight bits and a no-wait-state system design. The "no option" of each instruction used.

TABLE B-1. Average Execution Time

INSTRUCTION	NUMBER OF CYCLES
BROR	$42 + (107 + 44 \cdot (\text{width} - 2)) \cdot \text{height}$
BRORR	$44 + (107 + 44 \cdot (\text{width} - 2)) \cdot \text{height}$
BRAND	$45 + (111 + 44 \cdot (\text{width} - 2)) \cdot \text{height}$
BDOR	$48 + (61 + 25 \cdot (\text{width} - 2)) \cdot \text{height}$
	$\text{shift} = 0$
	$48 + (74 + 32 \cdot (\text{width} - 2)) \cdot \text{height}$
	$\text{shift} = 1 - 8$
BRSTCO	$66 + (170 + 60 \cdot (\text{width} - 2)) \cdot \text{height}$
	$\text{if shift} = 0, 16$
BITWT	$\text{if shift} = 1 - 8, 28$
	$35 + (19 + 12 \cdot \text{width}) \cdot \text{height}$
	$35 + (13 + 12 \cdot \text{width}) \cdot \text{height}$
EXTBLT	$35 + (13 + 12 \cdot \text{width}) \cdot \text{height}$
	(no pre-read)
MOVWRB,W	$16 + 7 \cdot R2$
MOVWPD	$16 + 8 \cdot R2$
TBIT	$27 \text{ per bit tested}$
SBITS	$\text{if R2} = 25$
	39
	else
	42
SRTP	$8 + (R4 \cdot R2)$
MOVSN	$59 + (14 \cdot R0) + (2 \cdot R0 \cdot 4)$
MOVSW	$59 + (14 \cdot R0) + (2 \cdot R0 \cdot 4)$
MOVSD	$59 + (10 \cdot R0) + (2 \cdot R0 \cdot 8)$

A.2.2 Interpreting the Table

To calculate the execution time for a given instruction, complete the formula and evaluate. For example, calculate the time for a 10-word wide, 50-line high BIT operation, the completed formula is:

$$42 + (107 + 44 \cdot (10 - 2)) \cdot 50$$

or

$$42 + (107 + 352) \cdot 50 = 22,992 \text{ clocks or } 1.53 \text{ msec at } 15.1 \text{ MHz, } 0 \text{ wait states.}$$

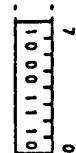
Appendix A: Instruction Formats (Continued)

Implied Immediate Encodings:



Trap (UND)

Format 17
Always



Format 18
Always

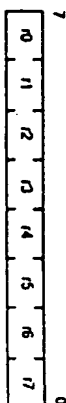


Trap (UND)

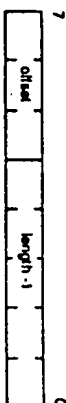
Format 19
Always



Register Mask, appended to SAVE, ENTER



Register Mask, appended to RESTORE, EXIT



Offset/Length Modifier appended to INNS, EXITS

Note 1: Opcode not defined; CPU treats the MOV, First operand has access class of read; second operand has access class of write; 1-field selects 32-bit or 64-bit data.

Note 2: Opcode not defined; CPU treats the ADD, First operand has access class of read; second operand has access class of read-modify-write; 1-field selects 32-bit or 64-bit data.

2.0 Architectural Description (Continued)

2.3.3 Instruction Set Summary

Table 2-2 presents a brief description of the NS32FX16 instruction set. The format column refers to the instruction format tables (Appendix A). The instruction column gives the instruction as coded in assembly language, and the Description column provides a short description of the function provided by that instruction. Further details of the exact operations performed by each instruction may be found in the *Series 32000 Instruction Set Reference Manual* and the *NS32CG16 Printer/Display Processor Programmer's Reference*.

Notations:

i = Integer length suffix: B = Byte

W = Word

D = Double Word

f = Floating Point length suffix: F = Standard Floating

L = Long Floating

MOVES

TABLE 2-2. NS32FX16 Instruction Set Summary

Format	Operation	Operands	Description
4	MOV	gen, gen	Move a value.
2	MOVCL	short, gen	Extend and move a signed 4-bit constant.
7	MOVWL	gen, gen, disp	Move multiple; disp bytes (1 to 16).
7	MOV/BW	gen, gen	Move with zero extension.
7	MOVZLD	gen, gen	Move with zero extension.
7	MOVXBLW	gen, gen	Move with sign extension.
7	MOVXBD	gen, gen	Move with sign extension.
4	MOVL	gen, gen	Move effective address.

INTEGER ARITHMETIC

Format	Operation	Operands	Description
4	ADD	gen, gen	Add.
2	ADDCL	short, gen	Add signed 4-bit constant.
4	ADDCL	gen, gen	Add with carry.
4	SUB	gen, gen	Subtract.
4	SUBCL	gen, gen	Subtract with carry (borrow).
6	NEG	gen, gen	Negate (2's complement).
6	ABS	gen, gen	Take absolute value.
7	MUL	gen, gen	Multiply.
7	QDQ	gen, gen	Divide, rounding toward zero.
7	REM	gen, gen	Remainder from QDQ.
7	DIV	gen, gen	Divide, rounding down.
7	MOD	gen, gen	Remainder from DIV (Modulus).
7	MULH	gen, gen	Multiply to extended integer.
7	DEH	gen, gen	Divide extended integer.

PACKED DECIMAL (BCD) ARITHMETIC

Format	Operation	Operands	Description
6	ADDP	gen, gen	Add packed.
6	SUBP	gen, gen	Subtract packed.

gen = General operand. Any addressing mode can be specified.

short = A 4-bit value encoded within the Basic Instruction (see Appendix A for encodings).

imm = Implied immediate operand. An 8-bit value appended after any addressing extensions.

disp = Displacement (addressing constant); 8, 16 or 32 bits. All three lengths legal.

reg = Any General Purpose Register: R0-R7.

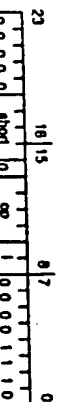
areg = Any Processor Register: SP, SB, FP.

MTBASE, MOD, PSR, US (bottom 8 PSR bits).

cond = Any condition code, encoded as a 4-bit field within the Basic Instruction (see Appendix A for encodings).

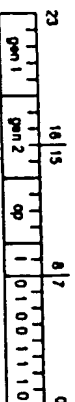
TABLE 2-2. NS32FX16 Instruction Set Summary (Continued)

INTEGER COMPARISON			
Format	Operation	Operands	Description
4	CMF ₇	gen,gen	Compare.
2	CMF ₁₆	short,gen	Compare to signed 4 bit constant.
7	CMF ₁₆	gen,gen,disp	Compare multiple; disp bytes (1 to 16).
LOGICAL AND BOOLEAN			
Format	Operation	Operands	Description
4	AND	gen,gen	Logical AND.
4	OR	gen,gen	Logical OR.
4	BIC	gen,gen	Clear selected bits.
4	XOR	gen,gen	Logical exclusive OR.
6	COM	gen,gen	Complement all bits.
6	NOT	gen,gen	Boolean complement; LSB only.
2	Scndf	gen	Save condition code (cond) as a Boolean variable of size i.
SHIFTS			
Format	Operation	Operands	Description
6	LSH	gen,gen	Logical shift, left or right.
6	ASH	gen,gen	Arithmetic shift, left or right.
6	ROf	gen,gen	Rotate, left or right.
BIT FIELDS			
Bit fields are values in memory that are not aligned to byte boundaries. Examples are PACKED arrays and records used in Pascal. "Extract" instructions read and align a bit field. "Insert" instructions write a bit field from an aligned source.			
Format	Operation	Operands	Description
8	EXT	reg,gen,gen,disp	Extract bit field (array oriented).
8	INS	reg,gen,gen,disp	Insert bit field (array oriented)
7	EXTS	gen,gen,imm,imm	Extract bit field (short form).
7	INSS	gen,gen,imm,imm	Insert bit field (short form).
8	CVIP	reg,gen,gen	Convert to bit field pointer.
ARRAYS			
Format	Operation	Operands	Description
8	CHECK	reg,gen,gen	Index bounds check.
8	INDEX	reg,gen,gen	Recursive indexing step for multiple-dimensional arrays.



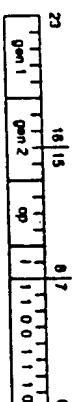
Format 5

MOV _S	-0000	IMWIT	-1000
CMF ₇	-0001	IMIS	-1001
SETCIG	-0010	IMAND	-1010
SKIPS	-0011	SHIPS	-1011
RBSTOD	-0100	BRFOR	-1100
EXTRLT	-0101	SMIS	-1101
BIOR	-0110	BBXOR	-1110
MOVAP	-0111		
No Operation on 1111			



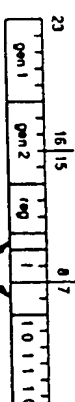
Format 6

ROT	-0000	NEG	-1000
ASH	-0001	NOT	-1001
COIT	-0010	Trap(UND)	-1010
CBITI	-0011	SUBP	-1011
Trap(UND)	-0100	ADS	-1100
LSH	-0101	COM	-1101
SOIT	-0110	IDIT	-1110
SOITI	-0111	ADOP	-1111



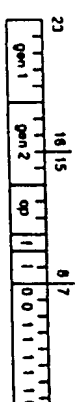
Format 7

MOV _M	-0000	MUL	-1000
CMF ₁₆	-0001	MEI	-1001
INSS	-0010	Trap(UND)	-1010
IXIS	-0011	QLO	-1011
MOV _{XIMW}	-0100	DIR	-1100
MOV _{ZIMW}	-0101	TELM	-1101
MOV _{ZID}	-0110	MCD	-1110
MOV _{XID}	-0111	DIV	-1111



Format 8

EXT	-0000	INDEX	-1000
CVIP	-0001	FFS	-1001
INS	-0010		
CHKCHK	-0011		
Trap(UND) on -110 and -111			



Format 9

MOV _I	-0000	ROUND	-1000
UTSR	-0001	TRAC	-1001
MOV _{FL}	-0010	SFSN	-1010
MOV _{FL}	-0011	FLOOR	-1011



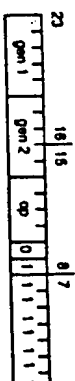
Format 10

Trap(UND) Always			
------------------	--	--	--



Format 11

ADDI	-0000	DVI	-1000
MOV _I	-0001	(Note 1)	-1001
CMF ₇	-0010	Trap(UND)	-1010
SUBI	-0011	Trap(UND)	-1011
NEGI	-0100	MULT	-1100
Trap(UND)	-0101	ABSI	-1101
Trap(UND)	-0110	Trap(UND)	-1110
Trap(UND)	-0111	Trap(UND)	-1111



Format 12

(Note 2)	-0000	(Note 2)	-1000
(Note 1)	-0001	(Note 1)	-1001
POLFI	-0010	Trap(UND)	-1010
DOFI	-0011	Trap(UND)	-1011
SCALBI	-0100	(Note 2)	-1100
LOGBI	-0101	(Note 1)	-1101
Trap(UND)	-0110	Trap(UND)	-1110
Trap(UND)	-0111	Trap(UND)	-1111



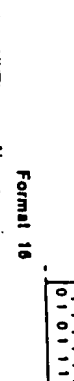
Format 13

Trap(UND) Always			
------------------	--	--	--



Format 14

Trap(UND) Always			
------------------	--	--	--



Format 15



Format 16

Appendix A: Instruction Formats

NOTATIONS

- Integer Type Field
 0 = 00 (Byte)
 1 = 01 (Word)
 2 = 11 (Double Word)
- Floating-Point Type Field
 F = 1 (Std. Floating: 32 bits)
 L = 0 (Long Floating: 64 bits)

Operation Code

Valid encoding shown with each format.
 gen 1, gen 2 = General Addressing Mode Field.
 See Section 2.3.2 for encodings.

General Purpose Register Number

- Condition Code Field
 0000 = Equal: Z = 1
 0001 = Not Equal: Z = 0
 0010 = Carry Set: C = 1
 0011 = Carry Clear: C = 0
 0100 = Higher: L = 1
 0101 = Lower or Same: L = 0
 0110 = Greater Than: N = 1
 0111 = Less or Equal: N = 0
 1000 = Flag Set: F = 1
 1001 = Flag Clear: F = 0
 1010 = Lower: L = 0 and Z = 0
 1011 = Higher or Same: L = 1 or Z = 1
 1100 = Less Than: N = 0 and Z = 0
 1101 = Greater or Equal: N = 1 or Z = 1
 1110 = (Unconditionally True)
 1111 = (Unconditionally False)

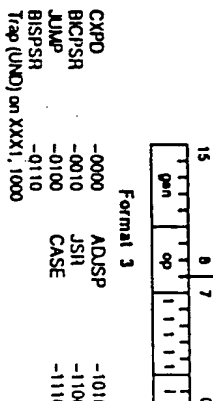
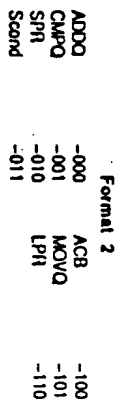
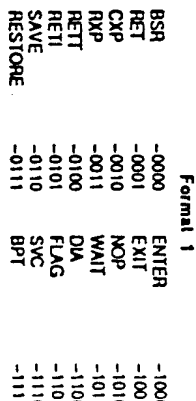
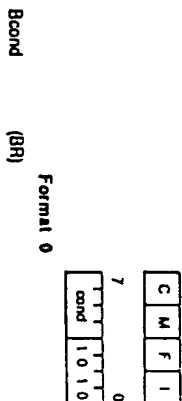
Short = Short immediate value. May contain quick: Signed 4-bit value. In MOVQ, ADDQ, CMPO, ACB.
 cond: Condition Code (above), in SCND, areg: CPU Dedicated Register, in LPR, SPR.

- 0000 = US
 0001 = 0111 = (Reserved)
 1000 = FP
 1001 = SP
 1010 = SB
 1011 = (Reserved)
 1100 = (Reserved)
 1101 = PSR
 1110 = MIBASE
 1111 = MOD

Options: in String Instructions

- T = Translated
 B = Backward
 U = None
 01: While Match
 11: Until Match

Configuration bits in SETCFG instruction



2.0 Architectural Description (Continued)

TABLE 2.2. NS32FX16 Instruction Set Summary (Continued)

STRINGS

- String instructions assign specific functions to the General Purpose Registers:
 R4 = Companion Value
 R2 = Transition Table Pointer
 R1 = String 2 Pointer
 R0 = String 1 Pointer
 R0 = Limit Count

Format	Operation	Operands
5	MOVSI	options
5	MOYST	options
5	CMPSI	options
5	SKPSI	options
5	SKPST	options

Description
 Move string 1 to string 2.
 Move string, translating bytes.
 Compare string 1 to string 2.
 Compare, translating string 1 bytes.
 Skip over string 1 entries.
 Skip, translating bytes for unilw/wh.

JUMPS AND LINKAGE

Format	Operation	Operands
3	JUMP	gen
0	BR	disp
0	BRCOND	disp
3	CASEI	gen
2	ACBI	short, gen, disp
3	JSR	gen
3	BSR	disp
3	CXP	disp
3	CRD	gen
1	SVC	disp
1	FLAG	disp
1	BPT	disp
1	ENTER	[reg list], disp
1	EXIT	[reg list]
1	RET	disp
1	RXP	disp
1	RETI	disp

Description
 Jump.
 Branch (PC Relative).
 Conditional branch.
 Multiway branch.
 Add 4-bit constant and branch if non-zero.
 Jump to subroutine.
 Branch to subroutine.
 Call external procedure.
 Call external procedure using descriptor.
 Supervisor call.
 Flag trap.
 Breakpoint trap.
 Save registers and allocate stack frame (Enter Procedure).
 Restore registers and reclaim stack frame (Exit Procedure).
 Return from subroutine.
 Return from external procedure call.
 Return from trap. (Privileged)
 Return from interrupt. (Privileged)

CPU REGISTER MANIPULATION

Format	Operation	Operands
1	SAVE	[reg list]
1	RESTORE	[reg list]
2	LPR	areg, gen
2	SPR	areg, gen
3	ADSPR	gen
3	BISPSR	gen
3	BICPSR	gen
5	SETCFG	[option list]

Description
 Save general purpose registers.
 Restore general purpose registers.
 Load dedicated register. (Privileged if PSR or INITBASE)
 Store dedicated register. (Privileged if PSR or INITBASE)
 Adjust stack pointer.
 Set selected bits in PSR. (Privileged if not Byte length)
 Clear selected bits in PSR. (Privileged if not Byte length)
 Set configuration register. (Privileged)

2.0 Architectural Description (Continued)

Table 2-2. NS32FX16 Instruction Set Summary (Continued)

FLOATING-POINT			
Format	Operation	Operands	Description
11	MOV	gen, gen	Move a floating-point value.
9	MOVLF	gen, gen	Move and shorten a long value to standard.
9	MOVFL	gen, gen	Move and lengthen a standard value to long.
9	MOVIL	gen, gen	Convert any integer to standard or long floating.
9	ROUND	gen, gen	Convert to integer by rounding.
9	TRUNC	gen, gen	Convert to integer by truncating, toward zero.
9	FLOOR	gen, gen	Convert to largest integer less than or equal to value.
11	ADD	gen, gen	Add.
11	SUB	gen, gen	Subtract.
11	MULT	gen, gen	Multiply.
11	DIV	gen, gen	Divide.
11	CMPL	gen, gen	Compare.
11	NEGL	gen, gen	Negate.
11	ABS	gen, gen	Take absolute value.
9	FSR	gen	Load FSR.
9	STFSR	gen	Store FSR.
12	POLY	gen, gen	Polynomial Step.
12	DOT	gen, gen	Dot Product.
12	SCALE	gen, gen	Binary Scale.
12	LOGBI	gen, gen	Binary Log.
MISCELLANEOUS			
1	NOP		No operation.
1	WAIT		Wait for interrupt.
1	DIA		Diagnose. Single-byte "Branch to Self" for hardware breakpointing. Not for use in programming.
GRAPHICS			
Format	Operation	Operands	Description
5	ROR	options*	Bit-aligned block transfer 'OR'.
5	IRAND	options	Bit-aligned block transfer 'AND'.
5	IROR	options	Bit-aligned block transfer fast 'OR'.
5	UIXOR	options	Bit-aligned block source to destination.
5	BSTOD	options	Bit-aligned word transfer.
5	BITWT	options	External bit-aligned block transfer.
5	EXTLT	options	Move multiple pattern.
5	MOVMP	options	Test bit string.
5	TRIS	options	Set bit string.
5	SRIS		Set bit perpendicular string.
5	SRIPS		
BITS			
Format	Operation	Operands	Description
4	TBIT	gen, gen	Test bit.
6	SBIT	gen, gen	Test and set bit.
6	SBITL	gen, gen	Test and set bit, interlocked.
6	CBIT	gen, gen	Test and clear bit.
6	CBITL	gen, gen	Test and clear bit, interlocked.
6	IBIT	gen, gen	Test and invert bit.
8	FFS	gen, gen	Find first set bit.

*Note: Options are controlled by bits of the instruction, PSR status bits, or dedicated register values.

4.0 Device Specifications (Continued)

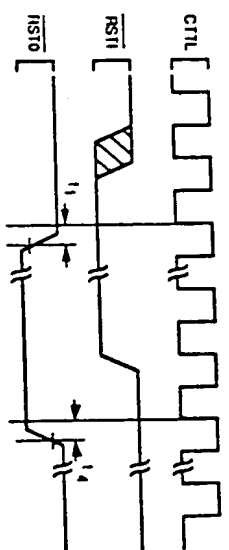


FIGURE 4-17. Non-Power-On Reset

Note 1: During Reset the \overline{RSTO} signal must be kept high.
Note 2: After \overline{RSTI} is deasserted the first bus cycle will be an instruction fetch at address zero.

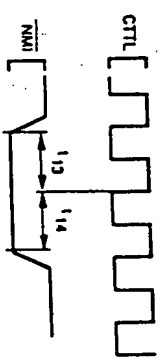


FIGURE 4-18. NMI Interrupt Signal Timing

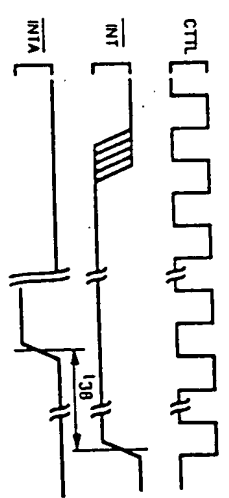


FIGURE 4-19. \overline{INT} Interrupt Signal Detection

Note 1: Once \overline{INT} is asserted, it must remain asserted until it is acknowledged.
Note 2: \overline{INTA} is the Interrupt Acknowledge bus cycle (not a CPU signal). Refer to Section 3.4.1 and Table 3.4.

4.0 Device Specifications (Continued)

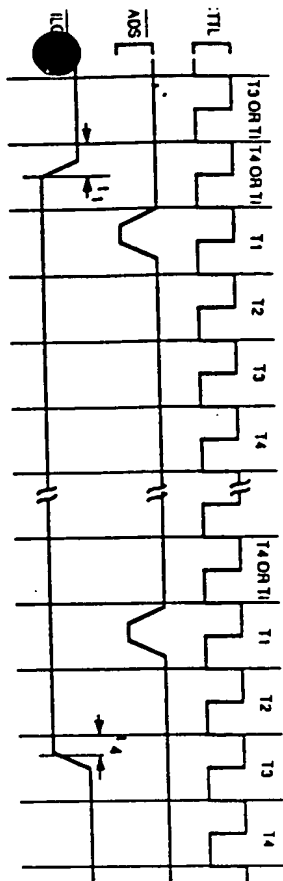


FIGURE 4-14. Interlocked Bus Cycle

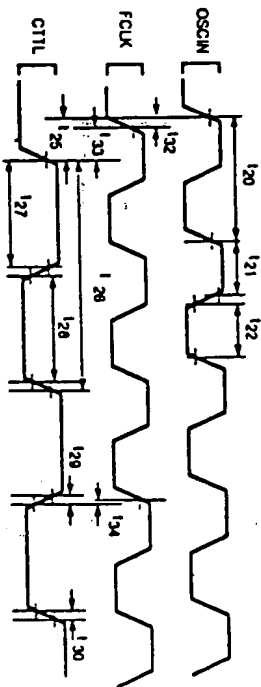


FIGURE 4-15. Clock Waveforms

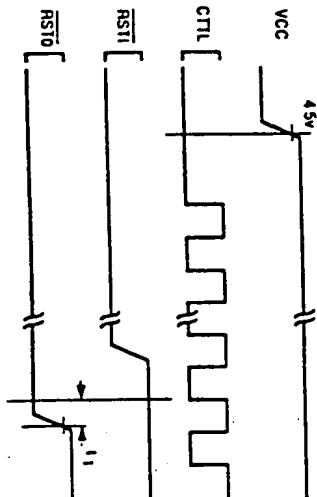


FIGURE 4-16. Power-On Reset

2.0 Architectural Description (Continued)

2.4 GRAPHICS SUPPORT

The following sections provide a brief description of the NS32FX16 graphics support capabilities. Basic discussions on frame buffer addressing, bilinear operations are also provided. More detailed information on the NS32FX16 graphics support instructions can be found in the NS32C616 Printer/Display Processor Programmer's Reference.

2.4.1 Frame Buffer Addressing

There are two basic addressing schemes for referencing pixels within the frame buffer: Linear and Cartesian (or x-y). Linear addressing associates a single number to each pixel representing the physical address of the corresponding bit in memory. Cartesian addressing associates two numbers to each pixel representing the x and y coordinates of the pixel relative to a point in the Cartesian space taken as the origin. The Cartesian space is generally defined as having the origin in the upper left. A movement to the right increases the x coordinate; a movement downward increases the y coordinate.

The correspondence between the location of a pixel in the Cartesian space and the physical (bit) address in memory is shown in Figure 2-10. The origin of the Cartesian space (x = 0, y = 0) corresponds to the bit address "ORC". Incrementing the x coordinate increments the bit address by one. Incrementing the y coordinate increments the bit address by an amount representing the warp (or pitch) of the Cartesian space. Thus, the linear address of a pixel at location (x, y) in the Cartesian space can be found by the following expression.

$$ADDR = ORC + Y * WARP + X$$

Warp is the distance (in bits) in the physical memory space between two vertically adjacent bits in the Cartesian space.

Example 1 below shows two NS32FX16 instruction sequences to set a single pixel given the x and y coordinates. Example 2 shows how to create a fat pixel by setting four adjacent bits in the Cartesian space.

Example 1: Set pixel at location (x, y)

Setup: R0 x coordinate
R1 y coordinate

Instruction Sequence 1:

```

MULD  WARP, R1      ; Y * WARP
ADDU   R0, R1        ; X + Y * WARP
SHLTD  R1, ORC        ; SET PIXEL

```

Instruction Sequence 2:

```

INDEXD R1, (WARP-1), R0 ; Y * WARP + X
SHLTD  R1, ORC          ; SET PIXEL

```

Example 2: Create fat pixel by setting bits at locations (x, y), (x + 1, y), (x, y + 1) and (x + 1, y + 1).

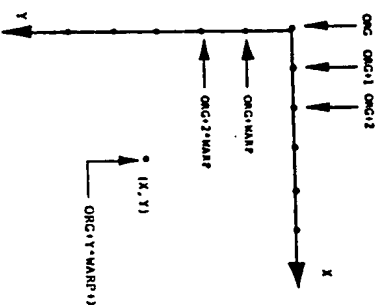
Setup: R0 x coordinate
R1 y coordinate

Instruction Sequence:

```

INDEXD R1, (WARP-1), R0 ; BIT ADDRESS
SHLTD  R1, ORC          ; SET FIRST PIXEL
ADDU   R1, R1           ; (X+1), Y
SHLTD  R1, ORC          ; SECOND PIXEL
ADDU   R1, ORC          ; (X, Y+1)
SHLTD  R1, ORC          ; THIRD PIXEL
ADDU   R1, R1           ; (X+1), Y+1
SHLTD  R1, ORC          ; LAST PIXEL

```



2.4.2.1 Frame Buffer Architecture

There are two basic types of frame buffer architectures: plane-oriented or pixel-oriented. BiBLIT takes advantage of the plane-oriented frame buffer architecture's attribute of multiple, adjacent pixels-per-word, facilitating the movement of large blocks of data. The source and destination starting addresses are expressed as pixel addresses. The width and height of the block to be moved are expressed in terms of pixels and scan lines. The source block may start and end at any bit position of any word, and the same applies for the destination block.

2.4.2.2 Bit Alignment

Before a logical operation can be performed between the source and the destination data, the source data must first be bit aligned to the destination data. In Figure 2-11 the source data needs to be shifted three bits to the right in order to align the first pixel (i.e., the pixel at the top left corner) in the source data block to the first pixel in the destination data block.

2.4.2.3 Block Boundaries and Destination Mask

Each BiBLIT destination scan line may start and end at any bit position in any data word. The neighboring bits (bits sharing the same word address with any words in the destination data block, but not a part of the BiBLIT rectangle) of the BiBLIT destination scan line must remain unchanged after the BiBLIT operation.

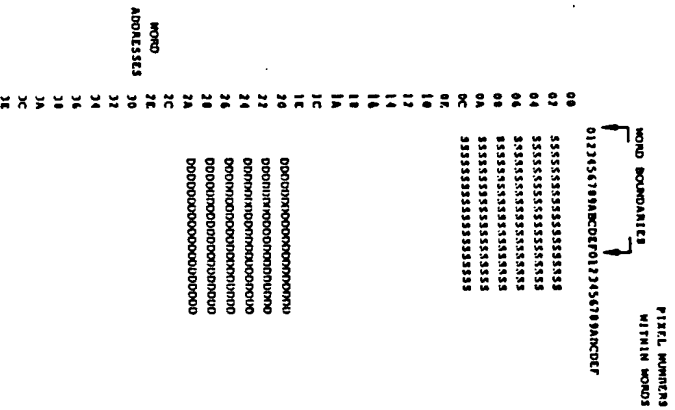


FIGURE 2-11. 32-Pixel Line Frame Buffer

Due to the plane-oriented frame buffer architecture, all memory operations must be word-aligned. In order to preserve the neighboring bits surrounding the BiBLIT destination block, both a left mask and a right mask are needed for all the leftmost and all the rightmost data words of the destination block. The left mask and the right mask both remain the same during a BiBLIT operation.

The following example illustrates the bit alignment requirements. In this example, the memory data path is 16 bits wide. Figure 2-11 shows a 32 pixel by 32 scan line frame buffer which is organized as a long bit stream which wraps around every two words (32 bits). The origin (top left corner) of the frame buffer starts from the lowest word in memory (word address 00 (hex)). Each word in the memory contains 16 bits. DO-D15. The least significant bit of a memory word, DO, is defined as the first displayed pixel in a word. In this example, BiBLIT addresses are expressed as pixel addresses relative to the origin of the frame buffer. The source block starting address is 021 (hex) (the second pixel in the 11th word). The destination block starting address is 204 (hex) (the 11th pixel in the 33rd word). The block width is 14 (hex), and the height is 06 (hex) (corresponding to 6 scan lines). The shift value is 3.

4.0 Device Specifications (continued)

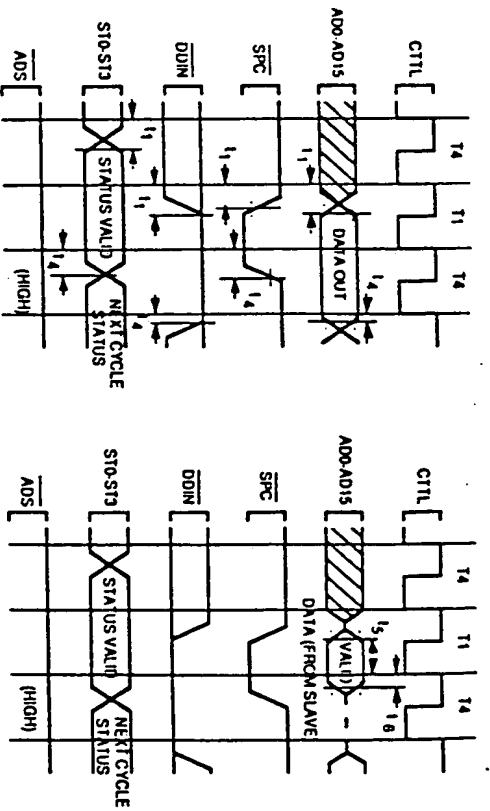


FIGURE 4-10. Slave Processor Write Timing

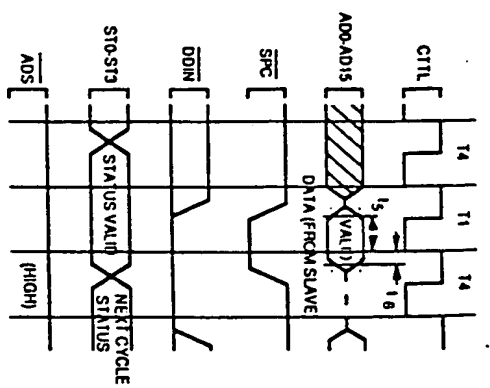


FIGURE 4-11. Slave Processor Read Timing

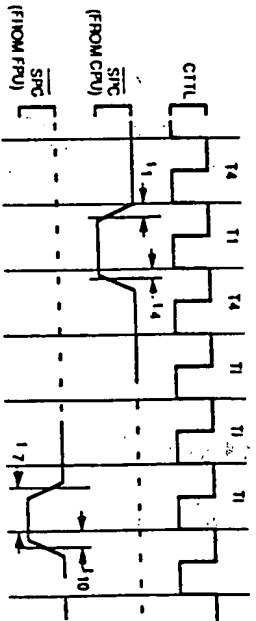


FIGURE 4-12. SPC Timing

There is a minimum one clock cycle between the SPC output asserted by the CPU and the SPC input from the FPU. After transferring the last operand to the FPU, the CPU turns OFF the output driver and holds SPC high with an internal SKI pulldown.

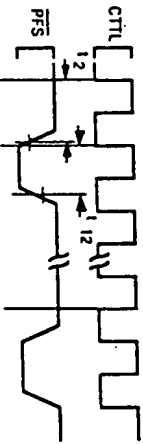
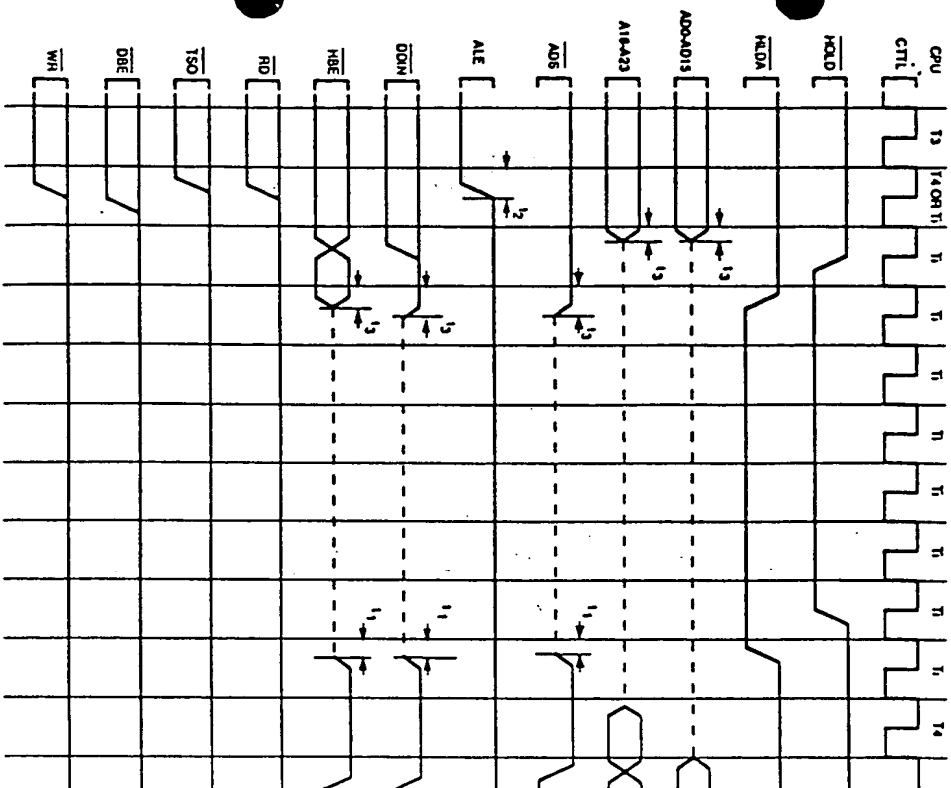
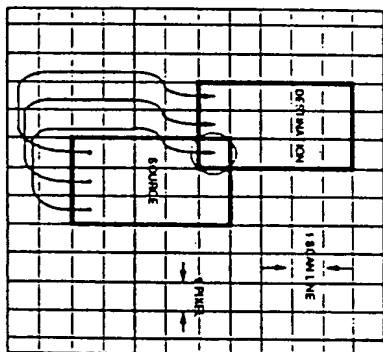


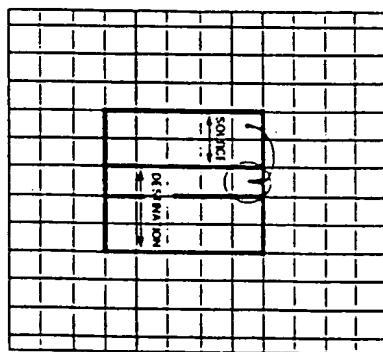
FIGURE 4-13. Relationship of PFS to Clock Cycles

FIGURE 4-9. HOLD Acknowledge (Bus Initially Idle)

2.0 Architectural Description (Continued)



(a)



(b)

FIGURE 2-12. Overlapping BILBT Blocks

The left mask and the right mask are 0000,1111,1111,1111 and 1111,1111,0000,0000 respectively.

Note 1: Zeros in either the left mask or the right mask indicate the destination bits which will not be modified.

Note 2: The BG (function) and EXBLT instructions use different setup parameters and techniques.

2.4.2.4 BILBT Directions

A BILBT operation moves a rectangular block of data in a frame buffer. The operation itself can be considered as a subroutine with two nested loops. The loops are preceded by setup operations. In the outer loop the source and destination starting addresses are calculated, and the test for completion is performed. In the inner loop the actual data movement for a single scan line takes place. The length of the inner loop is the number of (aligned) words spanned by each scan line. The length of the outer loop is equal to the height (number of scan lines) of the block to be moved. A skeleton of the subroutine representing the BILBT operation follows:

BILBT:

calculate BILBT setup parameters; (once per BILBT operation); such as width, height, bit misalignment (shift number), left, right masks, horizontal, vertical directions, etc.

OUTERLOOP:

calculate source, dest addresses; (once per scan line).

INNERLOOP:

move data (logical operation) and increment addresses; (once per word).

UNIL:

done horizontally.

UNIL:

done vertically.

Note:

In the NS32K X16 only the setup operations must be done by the programmer. The inner and outer loops are automatically executed by the BILBT instructions.

Each loop can be executed in one of two directions: the inner loop from left to right or right to left, the outer loop from top to bottom (down) or bottom to top (up).

The ability to move data starting from any corner of the BILBT rectangle is necessary to avoid destroying the BILBT source data as a result of destination writes when the source and destination are overlapped (i.e., when they share pixels). This situation is routinely encountered while panning or scrolling.

A determination of the correct execution directions of the BILBT must be performed whenever the source and destination rectangles overlap. Any overlap will result in the destruction of source data (from a destination write) if the correct vertical direction is not used. Horizontal BILBT direction is of concern only in certain cases of overlap, as will be explained below.

Figures 2-12(a) and (b) illustrate two cases of overlap. Here, the BILBT rectangles are three pixels wide by two scan lines high; they overlap by a single pixel in (a) and a single column of pixels in (b). For purposes of illustration, the BILBT is assumed to be carried out pixel-by-pixel. This convention does not affect the conclusions.

In Figure 2-12(a), if the BILBT is performed in the UP direction (bottom-to-top), one of the transfers of the bottom scan line of the source will write to the circled pixel of the destination. Due to the overlap, this pixel is also part of the uppermost scan line of the source rectangle. Thus, data needed later is destroyed. Therefore, this BILBT must be performed in the DOWN direction. Another example of this occurs any time the screen is moved in a purely vertical direction, as in scrolling text. It should be noted that, in both of these cases, the choice of horizontal BILBT direction may be made arbitrarily.

Figure 2-12(b) demonstrates a case in which the horizontal BILBT direction may not be chosen arbitrarily. This is an instance of purely horizontal movement of data (panning). Because the movement from source to destination involves data within the same scan line, the incorrect direction of movement will overwrite data which will be needed later. In this example, the correct direction is from right to left.

2.4.2.5 BBBLT Variations

The 'classical' definition of BBBLT, as described in "Smaliak-80 The Language and its Implementation" by Addele Goldberg and David Robson, provides for three operands: source, destination and mask/attribute. This third operand is commonly used in monochrome systems to incorporate a stipple pattern into an area. These stipple patterns provide the appearance of multiple shades of grey in single-bit-per-pixel systems, in a manner similar to the 'halftone' process used in printing.

Feature op1 Source op2 Destination + Destination

While the NS32C160 and the external DPU (if used) are essentially two-operand devices, three-operand BBBLT operations can be implemented quite flexibly and efficiently by performing the two operations serially.

2.4.3 Graphics Support Instructions

The NS32FX16 provides eleven instructions for supporting graphics oriented applications. These instructions are divided into three groups according to the operations they perform. General descriptions for each of them and the related formats are provided in the following sections.

2.4.3.1 BBBLT (Bit-aligned Block Transfer)

This group includes seven instructions. They are used to move characters and objects into the frame buffer which will be printed or displayed. One of the instructions works in conjunction with an external BBBLT Processing Unit (BPU) to maximize performance. The other six are executed by the NS32FX16.

Bit-aligned Block Transfer

Syntax: BB(function) Options

Setup: R0 base address, source data
R1 base address, destination data

R2 shift value
R3 height (in lines)
R4 first mask
R5 second mask
R6 source warp (adjusted)
R7 destination warp (adjusted)
q(CSF) width (in words)

Function: AND, OR, XOR, FOR, STOD

Options: IA Increasing Address (default option).

When IA is selected, scan lines are transferred in the increasing BIT/BYTE order.

DA Decreasing Address.

S True Source (default option).

.S Inverted Source.

These five instructions perform standard BBBLT operations between source and destination blocks. The operations available include the following:

BBAND: SRC AND DST
BBOR: SRC OR DST
BBXOR: SRC XOR DST
BBFOR: SRC FOR DST
BBSTOD: SRC TO DST

'src' and 'dst' stand for 'True Source' and 'Inverted Source' respectively. 'dst' stands for 'Destination'.

Note 1: For speed reasons, the BB instructions require the masks to be specified with respect to the source block. In Figure 2-11 masking was defined relative to the destination block.

Note 2: The options .S and DA are not available for the BBFOR instruction.

Note 3: BBFOR performs the same operation as BBOR with IA and S options.

Note 4: IA and DA are mutually exclusive and so are S and .S.

Note 5: This width is defined as the number of source data to read.

Note 6: An odd number of bytes can be specified for the source warp. However, word alignment of source scan lines will result in faster execution.

The horizontal and vertical directions of the BBBLT operations performed by the above instructions, with the exception of BBFOR, are both programmable. The horizontal direction is controlled by the IA and DA options. The vertical direction is controlled by the sign of the source and destination warps. Figure 2-13 and Table 2-3 show the format of the BB instructions and the encodings for the 'op' and 'r' field

23	18	15	9	7	0
0 0 0 0 0 0	D X	S 0	op	1	0 0 0 0 1 1 1 0

• 0 is set when the DA option is selected
• S is set when the .S option is selected
• X is set for BBAND, and it is clear for all other BB instructions

FIGURE 2-13. BB Instructions Format

TABLE 2-3. 'op' and 'r' Field Encodings

Instruction	Options	'op' Field	'r' Field
BBAND	Yes	1010	11
BBOR	Yes	0110	01
BBXOR	Yes	1110	01
BBFOR	No	1100	01
BBSTOD	Yes	0100	01

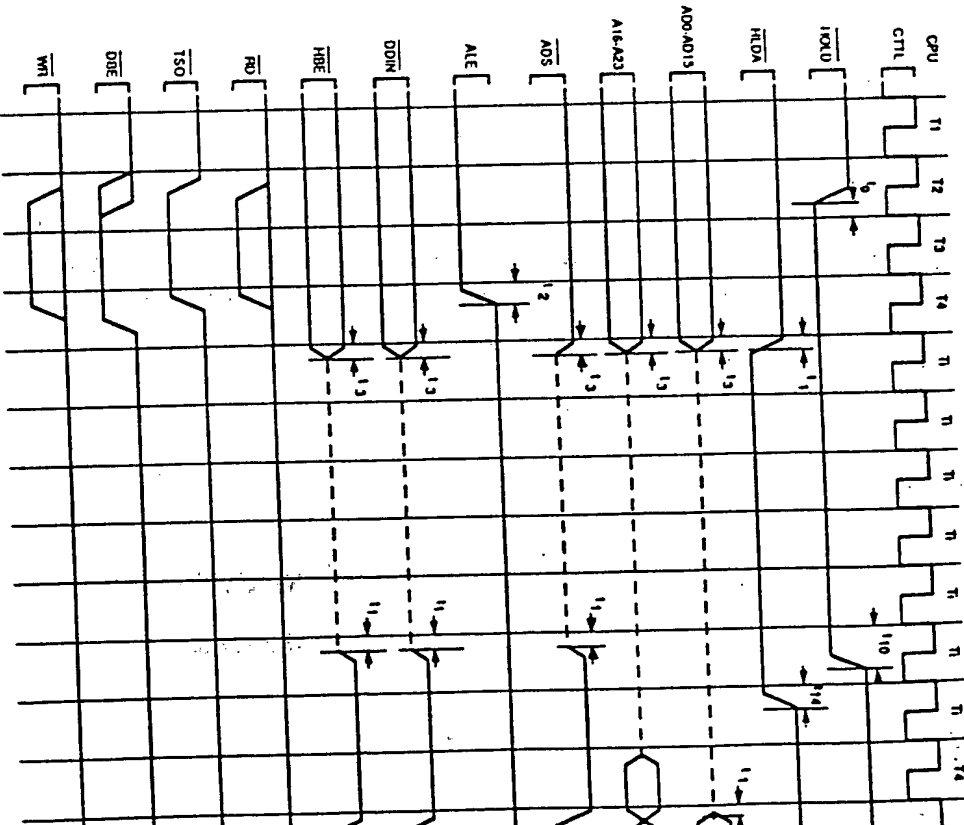


FIGURE 4-8. HOLD Acknowledge (Bus Initially Not Idle)

Note: When the bus is not idle, HOLD must be asserted before the rising edge of CTL or the timing starts that precedes state T4 in order for the request to be acknowledged.

This instruction shifts the bit for clear, adds a bit offset, and tests the bit for clear if "option" = 0 (and for set if "option" = 1). If clear (or set), the instruction increments to the next higher bit and tests for clear (or set). This testing for clear proceeds through memory until a set bit is found or until the maximum source bit offset or maximum run length value is reached. The total number of clear bits is stored in the destination as a run length value.

When TBITS finds a set bit and terminates, the bit offset is adjusted to reflect the current bit address. Offset is then ready for the next TBITS instruction with "option" = 0. After the instruction is executed, the F flag is set to the value of the bit previous to the bit currently being pointed to (i.e., the value of the bit on which the instruction completed execution). In the case of a starting bit offset exceeding the maximum bit offset (R1 \geq R4), the F flag is set if the option was 1 and clear if the option was 0. The L flag is set when the desired bit is found, or if the run length equaled the maximum run length value and the bit was not found. It is cleared otherwise. Figure 2-17 shows the TBITS instruction format.

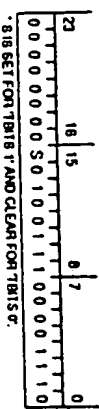


FIGURE 2-17. TBITS Instruction Format

Set Bit String

Syntax: SBITS

Setup: R0 base address of the destination

R1 starting bit offset (signed)

R2 number of bits to set (unsigned)

R3 address of string look-up table

Note: When the instruction terminates, the registers are retained unchanged.

SBITS sets a number of contiguous bits in memory to 1, and is typically used for data expansion operations. The instruction draws the number of ones specified by the value in R2, starting at the bit address provided by registers R0 and R1. In order to maximize speed and allow drawing of patterned lines, an external 1 Kbyte lookup table is used. The lookup table is specified in the NS32CG16 Printer/Display Processor Programmer's Reference Supplement.

When SBITS begins executing, it compares the value in R2 with 25. If the value in R2 is less than or equal to 25,

the F flag is cleared and the appropriate number of bits are set in memory. If R2 is greater than 25, the F flag is set and no other action is performed. This allows the software to use a faster algorithm to set longer strings of bits. Figure 2-18 shows the SBITS instruction format.



FIGURE 2-18. SBITS Instruction Format

Set Bit Perpendicular String

Syntax: SBITPS

Setup: R0 base address, destination (byte address)

R1 starting bit offset

R2 number of bits to set

R3 destination warp (signed value, in bits)

Note: When the instruction terminates, the R0 and R3 registers are retained unchanged. R1 becomes the final bit offset. R2 is zero.

The SBITPS can be used to set a string of bits in any direction. This allows a font to be expanded with a 90 or 270 degree rotation, as may be required in a printer application. SBITPS sets a string of bits starting at the bit address specified in registers R0 and R1. The number of bits in the string is specified in R2. After the first bit is set, the destination warp is added to the bit address and the next bit is set. The process is repeated until all the bits have been set. A negative raster warp offset value leads to a 90 degree rotation. A positive raster warp value leads to a 270 degree rotation. If the R3 value is = (space warp + 1 or -1), then the result is a 45 degree line. If the R3 value is +1 or -1, a horizontal line results.

SBITS and SBITPS allow expansion on any 90 degree angle, giving portrait, landscape and mirror images from one font. Figure 2-19 shows the SBITPS instruction format.

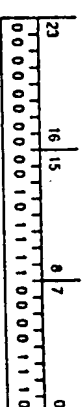
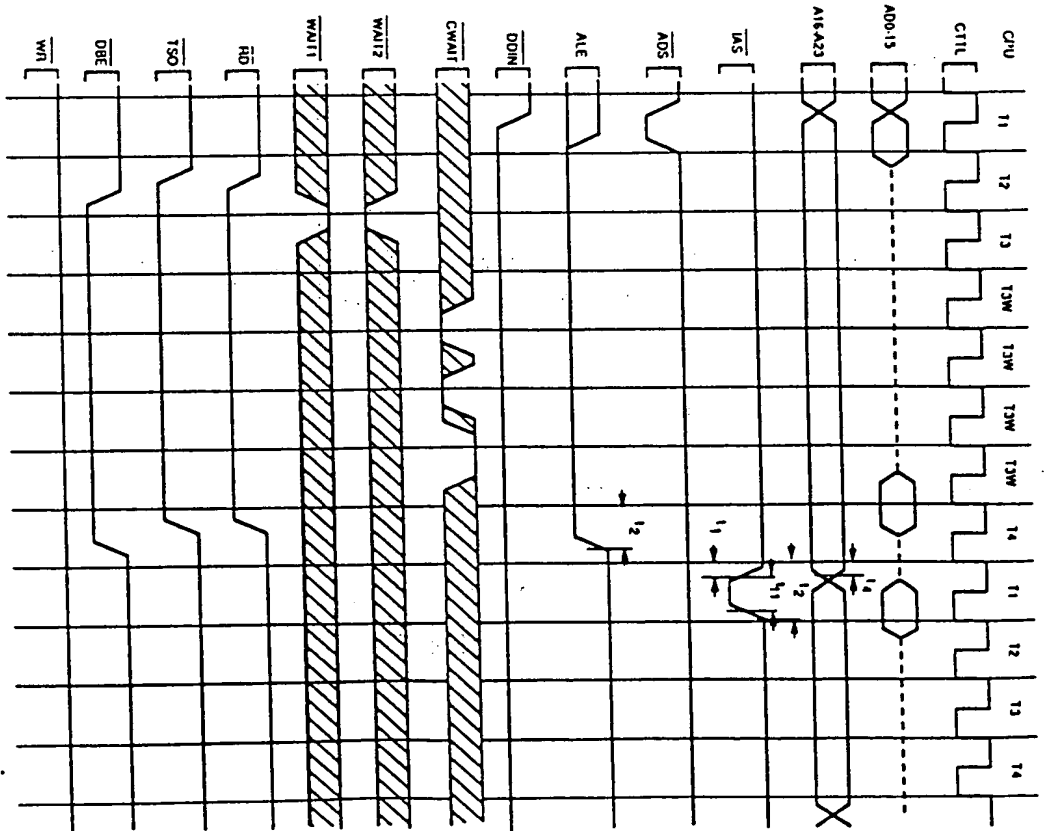


FIGURE 2-19. SBITPS Instruction Format



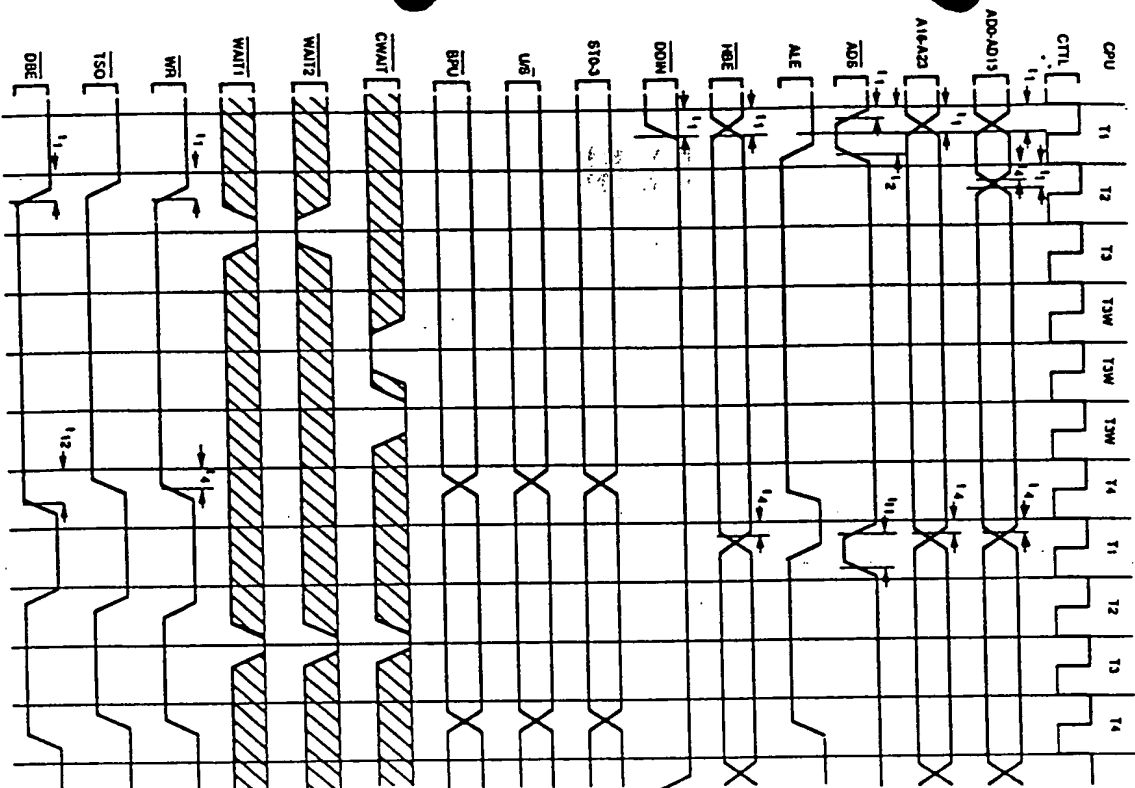


FIGURE 4-5. Write Cycle

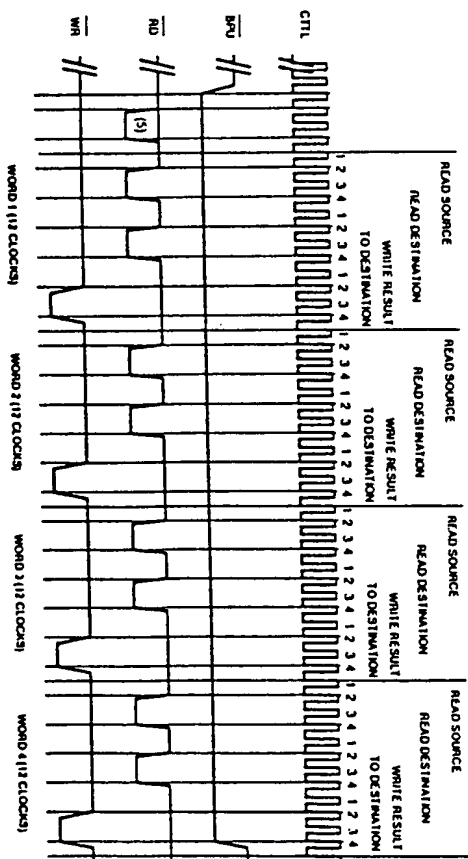


FIGURE 2-20. Bus Activity for a Simple BRBLT Operation

- Note 1: The example is for a block 4 words wide and 1 line high.
 Note 2: The sequence is common with all types of operations of the DP510DP511 BRU.
 Note 3: Data values, data values and number of bit planes do not affect the performance.
 Note 4: Zero wait states are assumed throughout the BRBLT operation.
 Note 5: The data read is performed when the BRU pipeline register needs to be preloaded.

2.4.3.3.1 Magnifying Compressed Data

Restoring data is just one application of the SDITS and SIIITPS instructions. Multiplying the "length" operand used by the SDITS and SIIITPS instructions causes the resulting pattern to be wider, or a multiple of "length". As the pattern of data is expanded, it can be magnified by 2x, 3x, 4x, ..., 10x and so on. This creates several sizes of the same style of character, or changes the size of a logo. A magnify in both dimensions X and Y can be accomplished by drawing a single line, then using the MOVs (Move String) or the UII instructions to duplicate the line, maintaining an equal aspect ratio.

More information on this subject is provided in the NS32CG16 Printer/Display Processor Programmer's Reference Supplement.

2.5 FAX ACCELERATOR MODULE

The FAX Accelerator Module (FAM) performs arithmetic operations on vectors of complex numbers. High performance is achieved by using a Hardware Multiplier Accumulator, an Address Generator for main memory operand accesses, and an on-chip RAM array.

The FAM executes complex arithmetic calculations on two vector operands. One vector is stored in the internal RAM array. The other vector is either organized as a circular buffer in the main memory or stored in the internal RAM array.

The FAM executes vector operations in a two stage pipeline. This allows for a significant performance enhancement as each operand fetch and execution on different vector elements are performed simultaneously rather than in a strictly sequential manner. The FAM can fetch up to two data elements at a time, using its address generator. The first operand is latched from the coefficient array whereas the second operand is from either external memory or from the coefficient array. While fetching the multiply and add vector element, the FAM performs the multiply and add operations on the previous vector element. Each complex multiply and accumulate operation requires two operand fetches, four multiplications and four additions. The FAM pipeline allows a maximal throughput of a complex multiply-accumulate operation in 8 clock cycles.

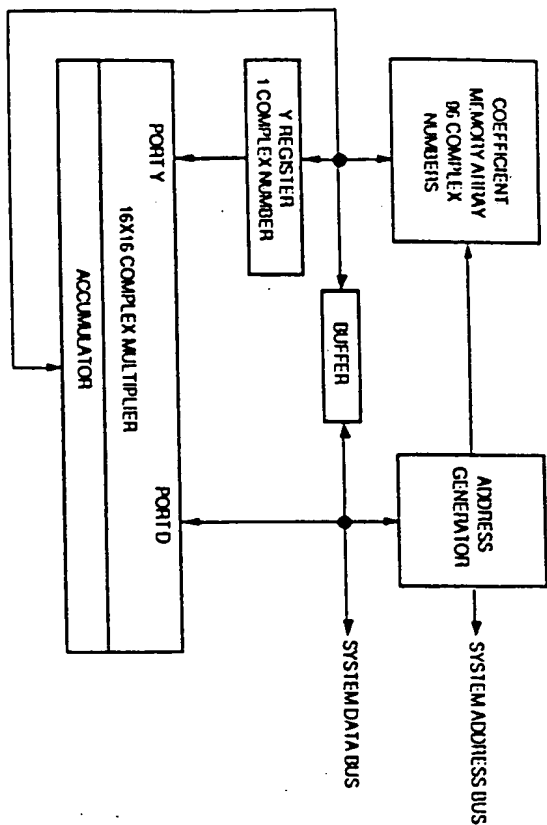


FIGURE 2-21. Fax Accelerator Module Block Diagram

2.5.1 FAM Operation

The following terms are used for the description of operations:

- C[i] Confident memory element, entry [i] can be selected by address generator or directly accessed by CPU.
- D[i] Data from external memory fetched using the address generator.
- Y Complex Multiplier input register.
- D[i]* The conjugate of D[i].
- A Complex Accumulator.

The FAX Accelerator Module can execute 6 basic commands:

- VCMAC Vector Complex Multiply Accumulate
- VCMAG Vector Complex Magnitude
- VCMAD Vector Complex Multiply Add
- VCMUL Vector Complex Multiply
- LOAD Write into C, Y, A, or CTL
- STORE Read from C, A, Y, ST or CTL

VCMAC, VCMAD and VCMUL commands use the following parameters:

- D vector starting address
- C vector starting address
- Vector length
- Control bits

VCMAG command uses only the last three operands.

2.5.1.1 Complex Number Representation

Complex numbers are organized as double words. Each double word contains two 16-bit 2's complement fractional integers. The less significant word contains the Real part of the number. The most significant word contains the Imaginary part of the number.

Complex vectors consist of arrays of complex numbers stored in consecutive addresses. Complex vectors MUST be aligned to double word boundary. Figure 2-22 illustrates the memory organization of vector D.

ADDRESS	CONTENTS
D	Re[D(0)]
D+2	Im[D(0)]
D+4	Re[D(1)]
D+6	Im[D(1)]
	.
D+4*n	Re[D(n)]
D+4*n+2	Im[D(n)]

FIGURE 2-22. Memory Organization of a Complex Vector

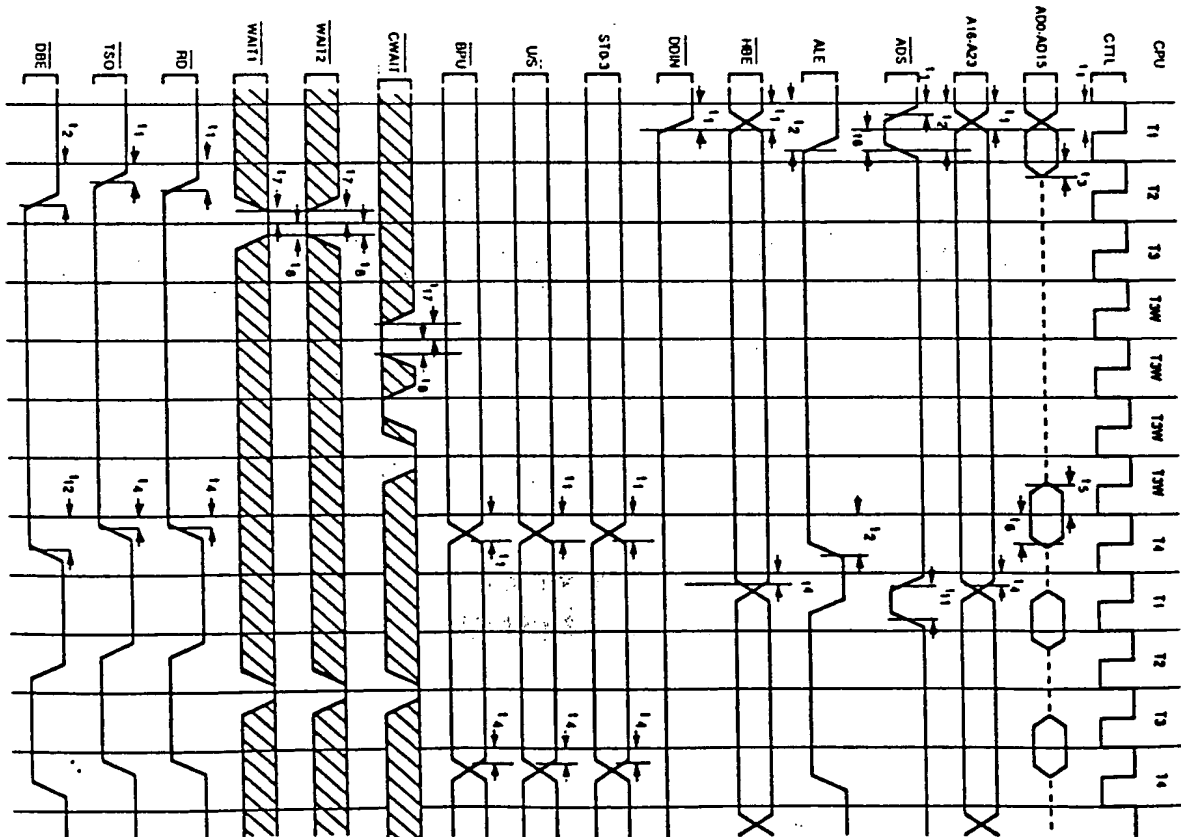


FIGURE 4-4. Read Cycle

4.0 Device Specifications (Continued)

4.4.2 Timing Tables

4.4.2.1 Output Signals: Internal Propagation Delays, NS32FX16-15, NS32FX16-20 and NS32FX16-25

Capacitive Load: CCTL -100pF, all other outputs -50pF

Symbol	Parameter	Reference	25 MHz	20 MHz	15 MHz	Units
t1	Output valid time	RE CCTL	12	13	14	ns
t2	Output valid time	HE CCTL	5126	5126	5126	ns
t3	Output hold time	RE CCTL	0	0	0	ns
t4	Output hold time	HE CCTL	0	0	0	ns
t5	Input setup time	RE CCTL	10	14	15	ns
t6	Input hold time	RE CCTL	2	2	2	ns
t7	Input setup time	HE CCTL	20	21	22	ns
t8	Input hold time	HE CCTL	2	2	2	ns
t9	Input setup time	RE CCTL	14	15	16	ns
t10	Input hold time	RE CCTL	2	2	2	ns
t11	Pulse width	0.8V	10	15	20	ns
t12	Output hold time	HE CCTL	5126	5126	5126	ns
t13	Input setup time	FE CCTL	12	14	15	ns
t14	Input hold time	FE CCTL	2	2	2	ns
t15	Input hold time	RE CCTL	2	2	2	ns
t16	Output valid to strobe inactive		10	10	10	ns
t17	Input setup time	RE CCTL	10	18	22	ns
t20	OSCN period	RE OSCIN	20	25	33	ns
t21	OSCN high time	4.2V	5120	5120	5120	ns
t211	OSCN low time	1.0V	5120	5120	5120	ns
t212	OSCN-CCTL delay	4.2V RE OSCIN	25	29	35	ns
t26	CCTL period	2.0V	40	50	66	ns
t27	CCTL high time	2.0V	5126	5126	5126	ns
t28	CCTL low time	0.8V	5126	5126	5126	ns
t29	CCTL fall time	RE CCTL	4	5	6	ns
t30	CCTL rise time	FE CCTL	4	5	6	ns
t38	NT signal hold after interrupt acknowledge		8	8	8	CCTL period
t32	OSCN to FCLK	4.2V RE OSCIN	15	20	25	ns
t33	FCLK to CCTL	RE FCLK	10	10	10	ns
t34	FCLK to CCTL	RE FCLK	10	10	10	ns

1. Not 100% tested

2.0 Architectural Description (Continued)

2.5.1.2 Mac Operation

The ALU of the FAX Accelerator Module contains a 16*16 multiplier and a 32-bit adder. Bits 15:30 (16 bits) of the result are rounded, and can be read by accessing the A register. If an overflow is detected during operation, the ST register OVF bit and either OPO or OP1 bits will be set to "1".

A 16-bit value is loaded into bits 15:30 of the Accumulator and the lower bits are set to "0". The value from bit 30 is copied into bit 31 for sign extension. Bit 14 is set to "1". An overflow is detected whenever the value of bit 30 is different from the value of bit 31.

2.5.1.3 Instruction Set

Each instruction of the FAM is controlled by two opcode bits (OPCO and OPC1), and two specifiers, COU and CLR. COU specifies whether or not the operand on port D of the multiplier needs to be conjugated prior to multiplication. The CLR bit is used to extend the instruction set. On VCMAC and VCMAG, CLR specifies whether or not the Accumulator has to be cleared at the beginning of the vector operation. On VCMAD, CLR is set to specify that the operation will ignore the value of CIL. In VCMUL, CLR is set to indicate that the value of DII is to be taken, instead of 1+DII. Table 2.4 is a summary of the various instruction sequences executed by the FAX Accelerator Module as a function of OPC1, OPC0, COU, and CLR bits in the CIL register.

The summation is done on N elements of the vector. All operands are complex numbers. Thus,

$$A \leftarrow \sum_{i=1}^N (C[i] \times D[i]) \text{ breaks down to:}$$

$$\text{Re}(A) = \sum_{i=1}^N (\text{Re}(C[i]) \times \text{Re}(D[i]) - \text{Im}(C[i]) \times \text{Im}(D[i]))$$

$$\text{Im}(A) = \sum_{i=1}^N (\text{Re}(C[i]) \times \text{Im}(D[i]) + \text{Im}(C[i]) \times \text{Re}(D[i]))$$

Note that the Accumulator (A), the multiplier input register (Y), the external data pointer (DPTR) and the coefficient pointer (CPT) registers are used as temporary registers during vector operations. The values previously stored in those registers are destroyed. If the contents of the Accumulator (A) register after a FAM operation is used as an initial value for the next FAM operation, it should be noted that the least significant bits of A (0-14) may contain a value other than zero.

2.5.1.4 Circular Buffers

The FAM accesses arrays of data in external memory using the DPTR as an address pointer. DS0 and DS1 bits of the CIL register control the size of the array. The FAM allows a convenient way of handling the data array as in a FIFO. Only the appropriate number of the least significant bits of the DPTR are incremented on each access. The upper bits remain constant. Table 2.5 shows which bits are incremented. The rest remain constant.

DS1	DS0	External Buffer Size(DM)	Constant Address bits	Incremented Address bits
0	0	8	A0, A5 A23	A1 A4
0	1	16	A0, A6 A23	A1 A5
1	0	32	A0, A7 A23	A1 A6
1	1	64	A0, A8 A23	A1 A7

TABLE 2.5 Circular Buffer Size

Instruction	OPC1	OPC0	CLR	COU	Operation
VCMAD	0	0	0	0	$C[i] \leftarrow C[i] + Y \times D[i]$
	0	0	1	0	$C[i] \leftarrow C[i] + Y \times D[i]^*$
	0	1	0	0	$C[i] \leftarrow Y \times D[i]$
	0	1	1	0	$C[i] \leftarrow Y \times D[i]^*$
VCMUL	0	1	0	0	$C[i] \leftarrow C[i] \times (1 + D[i])$
	0	1	0	1	$C[i] \leftarrow C[i] \times (1 + D[i]^*)$
	0	1	1	0	$C[i] \leftarrow C[i] \times D[i]$
	0	1	1	1	$C[i] \leftarrow C[i] \times D[i]^*$
VCMAC	1	0	0	0	$A \leftarrow A + \sum (C[i] \times D[i])$
	1	0	0	1	$A \leftarrow A + \sum (C[i] \times D[i]^*)$
	1	0	1	0	$A \leftarrow \sum (C[i] \times D[i])$
	1	0	1	1	$A \leftarrow \sum (C[i] \times D[i]^*)$
VCMAG	1	1	0	0	$A \leftarrow A + \sum (C[i] \times C[i])$
	1	1	0	1	$A \leftarrow A + \sum (C[i] \times C[i]^*)$
	1	1	1	0	$A \leftarrow \sum (C[i] \times C[i])$
	1	1	1	1	$A \leftarrow \sum (C[i] \times C[i]^*)$

TABLE 2.4 FAX Accelerator Instruction Set Summary

2.5.1.5 Performance Considerations

The FAX Accelerator Module is designed for optimal throughput in vector operations. Its two stage pipeline overlaps the execution of operand fetches and multiply-accumulate operations for different vector elements. The FAM can fetch up to two data elements at a time, using its address generator for main memory access and the coefficient array for the second operand. While fetching operands for one vector element, the FAM performs the multiplication and additions on the previous vector element. Each complex multiply and accumulate operation requires two operand fetches, four multiplications and four additions. The FAM pipeline allows a maximal throughput of a complex multiply accumulate operation in 8 clock cycles. See Section B4, FAX Accelerator Module Performance, for more details.

Access to the FAM registers while it is executing a vector operation are delayed (as if the CVALT input is active). When the FAM finishes the operation, access to the registers proceeds.

The FAM uses the full bandwidth of the external bus during VCMAD, VCMUL, or VCMAC operations. While executing the VCMAG instruction, the bus is free as no external operands are required. In this case the core CPU proceeds execution in parallel with the FAM operation.

During VCMAD, VCMAL or VCMAC operations, external HOLD requests will be granted at the end of each memory access. Note that interrupt requests cannot be acknowledged until the FAM finishes a vector operation.

2.5.2 FAM Registers and RAM Array

The FAM contains 7 registers and a 96 double-word RAM array. These registers and internal RAM can be accessed as memory-mapped I/O devices. Any reference to the registers and the RAM is done using the on-chip bus protocol. See Section 3.4.7.

All the registers, except for the Status Register (ST), are readable and writable. ST is read-only. Accessing the register memory locations should be a multiple of a byte-length. Word accesses must be on word boundary, and double-word accesses must be on double word boundary. Failing to do so will cause unpredictable results.

2.5.2.1 Coefficient RAM Array C[0]-C[95]

Each register in the coefficient array is 32-bits wide and holds one complex number. See Section 2.5.1.1.

Note that the RAM array is not limited to coefficient storage only. It can be used as a fast, zero-wait state on-chip memory for instructions and data storage.

However, the RAM can be used for instruction storage only if the instructions are loaded into this RAM using word-aligned accesses. This can be achieved by moving aligned double words from the external memory to the on-chip RAM. Data can also be stored in the on-chip RAM with one restriction: storing data in the on-chip RAM can be done only if all the data is written using aligned word or double-word accesses.

2.5.2.2 Multiplier Input Register Y

This 32-bit register holds one complex operand (see Section 2.5.1.1). The Y register is mapped into two consecutive words called Y0 and Y1.

2.5.2.3 Accumulator A

This 32-bit register holds one complex result (see Sec. 2.5.1.1). The A register is mapped into two consecutive words, also called A0 and A1.

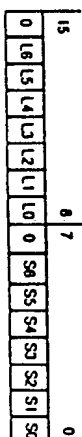
Internally, A0 and A1 are 32-bit registers, however, only bits 15:30 (16 bits) are accessible. The rest of the bits are used for a higher dynamic range on intermediate calculations.

2.5.2.4 Data Pointer DPTR

This is a 24-bit pointer to the beginning of the data vector in the main memory. In order to implement circular buffers, only the least significant bits of the DPTR pointer are incremented. When the end of a buffer is reached, the least significant bits of DPTR are loaded with zeros. The number of bits that are set to zero (which defines the size of the circular buffer) is controlled by CTL. The least-significant word of the DPTR is called DPTR0, and the most-significant byte is called DPTR1.

2.5.2.5 Coefficient Memory Vector Pointer CPTL

The CPTL register holds the address and length of the coefficient vector.



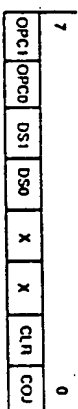
S0-S6 Start address of coefficient's vector (number of C.req).

L0-L6 Length of coefficient's vector (in double words).

Specifying 0 as the value of CPTL will cause an unpredictable result.

2.5.2.6 Control Register CTL

The CTL register controls the various modes of operation. For more details see Section 2.5.1.



OPC1-0 Operation code.

00 VCMAD Vector Complex Multiply Add
01 VCMAL Vector Complex Multiply
10 VCMAC Vector Complex Multiply Accumulate
11 VCMAG Vector Complex Magnitude

Register	Address	Length (bytes)	Direction
C[0]-C[95] reserved	FFFFD00 - FFFFD1F FFFD180 - FFFD1FF	96 x 4 = 384 640	R/W ...
Y	FFFD400	4	R/W
DPTR	FFFD408	4	R/W
CPTL	FFFD40C	2	R/W
CTL	FFFD40E	1	R/W
ST	FFFD410	1	R

TABLE 2.6 FAM Register Address Map

68-Pin PCC Package

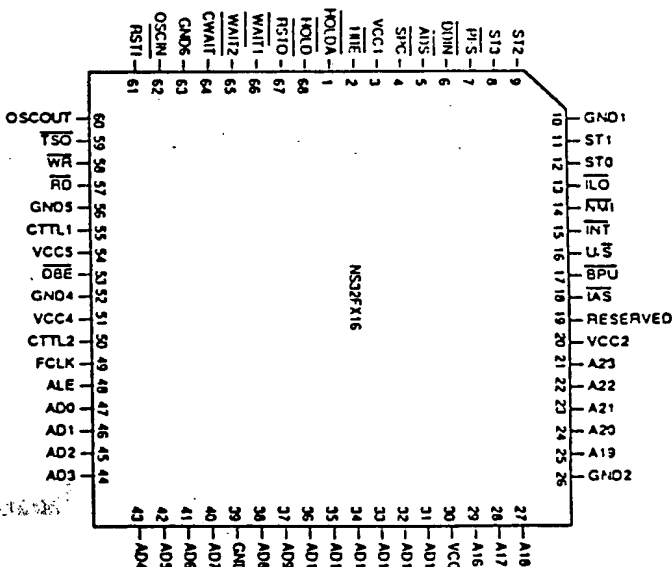


FIGURE 4-1. Connection Diagram

4.4 SWITCHING CHARACTERISTICS

4.4.1 Definitions

All the timing specifications given in this section refer to 0.8V or 2.0V on the rising or falling edges of CTL when the capacitive loading of CTL is 100 pF, unless specifically stated otherwise. The timing

specifications refer to 0.8 or 2.0V on the TTL output and input signals as illustrated in Figures 4-2 and 4-3, unless specifically stated otherwise.

ABBREVIATIONS:
R.E.—rising edge
L.E.—leading edge
T.E.—trailing edge
F.E.—falling edge

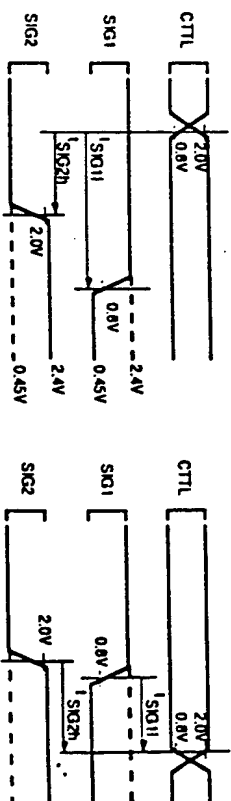


FIGURE 4-2. Timing Specifications Standard (Signal Valid After Clock Edge)

FIGURE 4-3. Timing Specification Standard (Signal Valid Before Clock Edge)

4.2 ABSOLUTE MAXIMUM RATINGS

Absolute maximum ratings indicate limits beyond which permanent damage may occur. Continuous operation at these limits is not intended; operation should be limited

4.3 ELECTRICAL CHARACTERISTICS: $I_A = 0$ C 10 0 0 0, $V_{CC} = 5$ V

Symbol	Parameter	Test Conditions	Min	Typ	Max	Unit
V _{IH}	High Level Input Voltage		2.0		V _{CC} + 0.5	V
V _{IL}	Low Level Input Voltage		-0.5		0.8	V
V _{XL}	OSCIN Input Low Voltage				0.5	V
V _{XH}	OSCIN Input High Voltage		4.5			V
V _{OH}	High Level Output Voltage	I _{OH} = -400 μ A	2.4			V
V _{OL}	Low Level Output Voltage	I _{OL} = 4 mA			0.45	V
I _I	Input Load Current	V _{CC} = 5.5V, V _{SS} = 0V, V _{IN} = 0.55V(1)	-20		20	μ A
I _L	Leakage Current	V _{CC} = 5.5V, V _{SS} = 0V, V _{OUT} = 0.4, 5.5V	-20		20	μ A
I _{IL} S	SPEC Input Current (twi)	V _{IN} = 0.4V, SPEC in Input Mode				mA
I _{TH} L	C11L Input Current (low)				4.0	mA
I _{CC}	Supply Current	25 MHz, T _A = 25°C, I _{OUT} = 0(2)		170		mA

Note: Applying voltage beyond that level might overstress the part. Applying voltage on logic pins beyond that level might cause latchup to the product.

As reported, output voltage is $-0.5V$ to $+7V$ (respected to GND)

Note 2: 1 CC is affected by the clock scaling factor selected by the C and M bits in the C/D register; see Section 3.2.1.

2.0 Architectural Description (Continued)

8 double-words

64 double-words 11

details, see section 2.5.1.3.

conjugated prior to multiplication.

The ST register holds the status of the last vector operation.

7						0
OVF	X	X	X	X	X	OP1 OP0

Q10 **Overflow occurred on calculation of A0.**

The ST register is cleared to 0 in the following cases

— upon reset.

3.0 Functional Description

The NS32FX16 requires a single 5-Volt power supply applied on 5 pins: VCC1-VCC5.

Grounding connections are made on 6 pins: GND1, GND6.

point. Daisy chained connections should be avoided.

VCC and ground. They should be connected to VCC as close as possible in the NS32F16

The NS32FX16 provides an internal oscillator that interacts with an external clock source through two signals: OSCIN and OSCOUT.

Either an external signal-phase clock signal or a crystal can be used as the clock source. If a single-phase clock source is used, only the connection on OSCIN is required; OSCOUT should be left unconnected or loaded with no more than 5 pF of stray capacitance. The voltage level requirements specified in Section 4.3 must also be met for proper operation.

When operation with a crystal is desired, special care should be taken to minimize stray capacitances and inductances. The crystal, as well as the external IIC components, should be placed in close proximity to the OSCIN and OSCOUT pins to keep the printed circuit trace lengths to an absolute minimum.

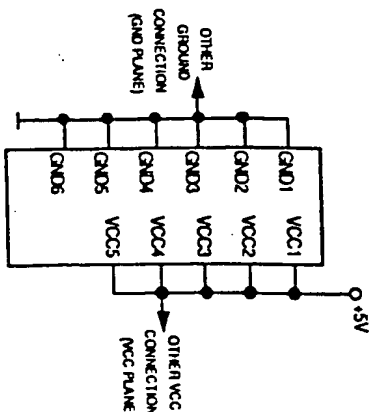


FIGURE 3.1. Power and Ground Connections

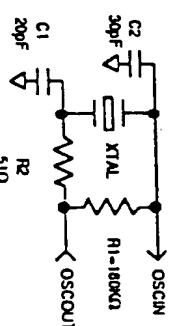


FIGURE 3-2(a). Crystal Interconnections - 30 MHz

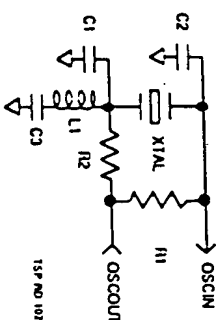


FIGURE 3-2(b). Crystal Interconnections - 40 MHz, 50 MHz

RCL Component Value						
Frequency (MHz)	R1 (k Ω)	C1 (pF)	C2 (pF)	C3 (pF)	L1 (μ H)	R2 (Ω)
40	150	20	20	200	1	51
50	150	20	20	200	0.6	51

TABLE 3-1. External Oscillator Specifications

Crystal Characteristics

Type	At-Cut
Tolerance	0.005% at 25°C
Stability	0.01% from 0°C to 70°C
Resonance	30 MHz - Fundamental (parallel)
Maximum Shunt Capacitance	7 pF
Maximum Series Resistance	50 Ω

3.2.1 Power Save Mode

The NS32FX16 provides a power save feature that can be used to significantly reduce the power consumption at times when the computational demand decreases. The device uses the clock signal at the OSCIN pin to derive the internal clock as well as the external signals CCTL and FCLK. The frequency is affected by the clock scaling factor. Scaling factors of 1, 2, 4 or 8 can be selected by properly setting the C and M bits in the CFG register. The power save mode should not be used to reduce the CCTL clock frequency below the minimum frequency required by the CPU (1 MHz).

Upon reset, both C and M are set to zero; thus, maximum clock rate is selected.

Due to the fact that the C and M bits are programmed by the SETCFG instruction, the power save feature can only be controlled by programs running in supervisor mode.

The following table shows the C and M bit settings for the various scaling factors, and the resulting supply current for a crystal frequency of 50 MHz.

Clock Scaling Factor vs Supply Current

C	M	Scaling Factor	CPU Clock Frequency	Typical I _{cc} at 5V
0	0	1	25 MHz	170 mA
0	1	2	12.5 MHz	91 mA
1	0	4	6.25 MHz	50 mA
1	1	8	3.13 MHz	30 mA

3.3 RESETTING

The RSTI input pin is used to reset the NS32FX16. The CPU samples RSTI on the falling edge of CCTL.

Whenever a low level is detected, the CPU responds

immediately. Any instruction being executed is terminated; any results that have not yet been written to memory are discarded; and any pending interrupts and traps are eliminated. The internal latch for the edge-sensitive NMI signal is cleared.

On application of power, RSTI must be held low for at least 50 μs after VCC is stable. This is to ensure that all on-chip voltages are completely stable before operation. Whenever a Reset is applied, it must also remain active for not less than 64 CCTL cycles. See Figures 3-4 and 3-5.

While in the Reset state, the CPU drives the signals AD5, RD, WR, DME, TSO, BDI, and (X)IN inactive. ADO, AD15, A16-A23 and SPC are floated, and the state of all other output signals is undefined.

The internal CPU clock and CCTL run at half the frequency of the signal on the OSCIN pin. FCLK runs at the same frequency as OSCIN.

The HOLD signal must be kept inactive. After the RSTI signal is driven high, the CPU will stay in the reset condition for approximately 8 clock cycles and then it will begin execution at address 0.

The PSR is reset to 0. The CFG C and M bits are reset to 0. RMI is enabled to allow Non-Maskable Interrupts. The following conditions are present after reset due to the PSR being reset to 0:

Tracing is disabled.

Supervisor mode is enabled.

Supervisor stack space is used when the TOS addressing mode is indicated.

No trace traps are pending.

Only RMI is enabled; IRI is not enabled.

IRU is inactive high.

The Clock Scaling Factor is set to 1; refer to Section 3.2.1.

Note that vector/non-vector interrupts have not been selected. While interrupts are disabled, a SETCFG [I] instruction must be executed to declare the presence of the NS32202 if vectored interrupts are desired. If non-vectored interrupts are required, a SETCFG without the [I] must be executed.

The presence/absence of the NS32081 or NS32381 has also not been declared. If there is a Floating-Point Unit, a SETCFG [F] instruction must be executed. If there is no floating-point unit, a SETCFG without the [F] must be executed.

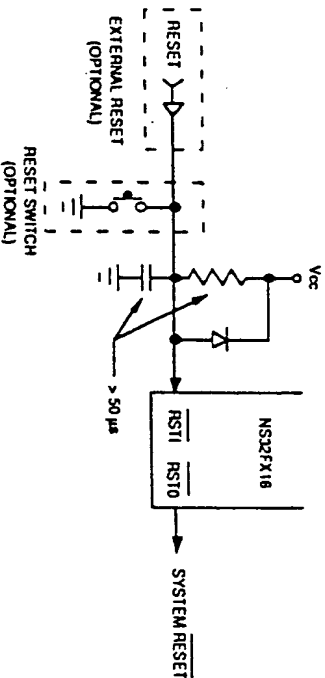


FIGURE 3-3. Recommended Reset Connections

STO-ST3

Bus Status
Bus cycle status code. ST0 is the least significant bit. Encodings are:

- 0000 — Idle: CPU inactive on bus.
- 0001 — Idle: WAIT instruction.
- 0010 — FAW Data Transfer.
- 0011 — Idle: Waiting for Slave.
- 0100 — Interrupt Acknowledge Master.
- 0101 — Interrupt Acknowledge, Cascaded.

- 0110 — End of Interrupt, Master.
- 0111 — End of Interrupt, Cascaded.
- 1000 — Sequential Instruction Fetch.
- 1001 — Non-Sequential Instruction Fetch.

BDI

- 1010 — Data Transfer.
- 1011 — Read Read-Modify-Write Operand.
- 1100 — Read for Effective Address.
- 1101 — Transfer Slave Operand.
- 1110 — Read Slave Status Word.
- 1111 — Broadcast Slave ID.

SPC

Timing State Output.
The falling edge of TSO identifies the beginning of state T2 of a bus cycle. The rising edge identifies the beginning of state T4.

U/S

User/Supervisor.
User or Supervisor Mode status. High indicates User Mode; low indicates Supervisor Mode.

WR

Write Strobe.
Activated during CPU or DMA write cycles to enable writing of data to memory or peripherals.

4.1.4 Input-Output Signals

ADO-AD15

Address/Data Bus.
Multiplexed Address/Data Information. Bit 0 is the least significant bit of each

AD5

Address Strobe.

Signals the beginning of a bus cycle and can be used for controlling the address latches. During HLD_A asserted, the signal becomes an input and the C monitors it to detect the beginning of external DMA cycle and generate relevant strobe signals. When external DMA Controller is used, AI should be pulled up to VCC through 10K resistor.

Data Direction.

Status signal indicating the direction the data transfer during a bus cycle. During HLD_A asserted, this signal becomes an input and determines activation of RD or WR.

Slave Processor Control.

Used by the CPU as the data strobe output for slave processor transfer used by a slave processor acknowledge completion of a slave instruction.

4.0 Device Specifications

4.1 PIN DESCRIPTIONS

4.1.1 Supplies
VCC1-5 Power
 +5V positive supplies
GND1-6 Ground

4.1.2 Input Signals
CWAIT Continuous Wait.
 Causes the CPU to insert continuous wait states if sampled low at the end of T2 and each following T3 or T3W state. WAIT1-WAIT2 inputs are sampled by the CPU during T3 or T3W if CWAIT asserted (low) and the corresponding wait-state counter is initialized. The wait states due to WAIT1-WAIT2 (if any) are added only after CWAIT is removed (becomes high). See Section 3.4.3.

HOLD

Hold Request.
 When active, causes the CPU to release the bus for DMA or multiprocessing purposes. See Section 3.5.

Note: If the HOLD signal is generated asynchronously, an setup and hold time may be required. In this case, it is recommended to synchronize it with CTT1 to minimize the possibility of metastable states.

The CPU provides only one synchronization stage to maintain the HOLD latency. This is to avoid speed degradation. In case of heavy HOLD activity (e.g., DMA controller cycles interleaved with CPU cycles).

Interrupt.
 A low level on this pin requests a maskable interrupt. INT must be asserted until the interrupt is acknowledged.

Non-Maskable Interrupt.
 A high-to-low transition on this signal requests a non-maskable interrupt.

Crystal/External Clock Input.
 Input from a crystal or an external clock source.

Reset Input.
 Asynchronous signal (Schmitt triggered) used to generate a CPU reset.

Wait State Inputs.

Wait1 Binary weighted inputs, allowing from zero to three wait states to be specified. The WAIT1-WAIT2 value is sampled by the CPU at the end of T2 (if CWAIT not asserted) or at the end of the last T3 or T3W state in which CWAIT is asserted. See Section 3.4.3.

4.1.3 Output Signals

A16-A23 High-Order Address Bits.
 These are the most significant 8 bits of the memory address bus.

ALE Address Latch Enable.
 Controls address latches.

BPU BPU Cycle.

Activated (low) during a bus cycle to enable an external BPU/T processing unit. The EXIBLT instruction activates this signal.

CTTL1-2 System Clock.
 CTTL1 and CTTL2 should be connected together externally.

DBE Data Buffers Enable.
 Used to control external data buffers. It is active when the data buffers are to be enabled.

FCLK Fast Clock.

This clock is derived from the clock waveform on OSCIN. Its frequency is either the same as OSCIN or is lower, depending upon the scale factor programmed into the CFG register. See Section 3.2.1.

HBE High Byte Enable.
 Status signal used to enable data transfers on the most significant byte of the data bus.

HOLDA

Hold Acknowledge.
 Activated by the CPU in response to the HOLD input to indicate that the CPU has released the bus.

Internal Address Strobe.

Signals the beginning of an on-chip bus cycle. IAS is a status signal used for debugging and tracing.

Interlocked Operation.
 When active (low), indicates that an interlocked operation is being executed.

Crystal Output.

This line is used as the return path for the crystal (if used). It must be left open when an external clock source is used to drive OSCIN.

Program Flow Status.
 A pulse on this line indicates the beginning of execution of an instruction.

Read Strobe.

Activated during CPU or DMA read cycles to enable reading of data from memory or peripherals.

Reset Output.
 Asserted (low) when RST1 is low, initiating a system reset.

3.0 Functional Description (Continued)

In general, a SETCFG instruction must be executed in the reset routine. In order to properly configure the CPU, the options should be combined and executed in a single instruction. For example, to declare vectored interrupts, a Floating-Point unit installed, and full CPU clock rate, execute a SETCFG [F, I] instruction. To declare non-vectored interrupts, no FPU, and full CPU clock rate, execute a SETCFG [I] instruction.

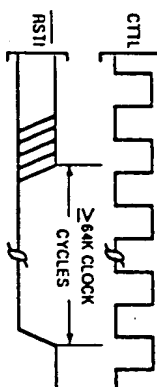


FIGURE 3-4. General Reset Timing

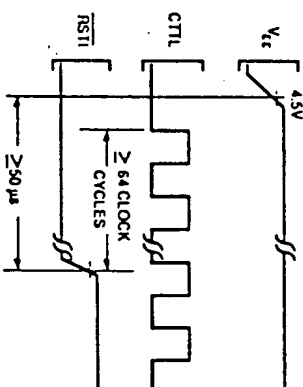


FIGURE 3-5. Power-on Reset Requirements

3.4 BUS CYCLES

The CPU will perform a bus cycle for one of the following reasons:

- 1) To write or read data, to or from memory or peripheral devices. Peripheral input and output are memory-mapped in National's Embedded System Processor family.
- 2) To fetch instructions into the eight-byte instruction queue. This happens whenever the bus would otherwise be idle and the queue is not already full.
- 3) To acknowledge an interrupt and allow external circuitry to provide a vector number, or to acknowledge completion of an interrupt service routine.
- 4) To transfer information to or from a Slave Processor.
- 5) To indicate an internal bus cycle (e.g., read of an on-chip FAM control register).

In terms of bus timing, cases 1 through 3 above are identical. For timing specifications, see Section 4. The only external difference between them is the four-bit code placed on the Bus Status pins (S10-S13).

Slave Processor cycles differ in that separate control signals are applied and there is no address involved (Section 3.4.9).

When using an external DMA channel, NS32FX16 still generates bus control signals if the DMA controller provides inputs that indicate the beginning of the DMA cycle (ADS) and the cycle type (DUN). However, the address is generated in this case by the external DMA Controller.

Case 5 does not represent an active bus cycle (ADS not asserted). Instead, a special address strobe, IAS, is asserted. The purpose of these cycles is to allow a debug or trace device (e.g., ISE) to track bus transactions inside the chip.

3.4.1 Bus Status

The NS32FX16 CPU presents four bits of Bus Status information on pins S10-S13. The various combinations on those pins indicate why the CPU is performing a bus cycle, or, if it is idle on the bus, then why it is idle.

The Bus Status pins are interpreted as a four-bit value, with S10 the least significant bit. Their values decode as follows:

- 0000 — The bus is idle because the CPU does not need to perform a bus access.
- 0001 — The bus is idle because the CPU is executing the WAIT instruction.
- 0010 — FAM Data Transfer.
- 0011 — The bus is idle because the CPU is waiting for a Slave Processor to complete an instruction.
- 0100 — Interrupt Acknowledge, Master.
 The CPU is performing a read cycle to acknowledge an interrupt request. See Section 3.4.6.
- 0101 — Interrupt Acknowledge, Cascaded.
 The CPU is reading an interrupt vector to acknowledge a maskable interrupt request from a Cascaded Interrupt Control Unit.
- 0110 — End of Interrupt, Master.
 The CPU is performing a read cycle to indicate that it is executing a Return from Interrupt (RETI) instruction at the completion of an interrupt's service procedure.
- 0111 — End of Interrupt, Cascaded.
 The CPU is performing a read cycle from a Cascaded Interrupt Control Unit to indicate that it is executing a Return from Interrupt (RETI) instruction at the completion of an interrupt's service procedure.
- 1000 — Sequential Instruction Fetch.
 The CPU is reading the next sequential word from the instruction stream into the instruction queue. It will do so whenever the bus would otherwise be idle and the queue is not already full.
- 1001 — Non-Sequential Instruction Fetch.
 The CPU is performing the first fetch of instruction code after the Instruction Queue

is purged. This will occur as a result of any jump or branch, any interrupt or trap, or execution of certain instructions.

1010 — Data Transfer.
The CPU is reading or writing an operand of an instruction.

1011 — Read RMW Operand.
The CPU is reading an operand which will subsequently be modified and rewritten. The write cycle of RMW will have a "write" status.

1100 — Read for Effective Address Calculation.
The CPU is reading information from memory in order to determine the Effective Address of an operand. This will occur whenever an instruction uses the Memory Relative or External addressing mode.

1101 — Transfer Slave Processor Operand.
The CPU is either transferring an instruction operand to or from a Slave Processor, or it is issuing the Operation Word of a Slave Processor instruction. See Section 3.4.9.2.

1110 — Read Slave Processor Status.
The CPU is reading a Status Word from a Slave Processor after the Slave Processor has signalled completion of an instruction.

1111 — Broadcast Slave ID.
The CPU is initiating the execution of a Slave Processor instruction by transferring the first byte of the instruction, which represents the slave processor identification.

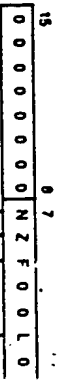
3.8.2 Floating-Point Instructions

Table 3-5 gives the protocols followed for each Floating-Point instruction. The instructions are referenced by their mnemonics. For the bit encodings of each instruction, see Appendix A.

The Operand class columns give the Access Class for each operand, defining how the addressing modes are interpreted (see *Series 32000 Instruction Set Reference Manual*).

The Operand issued columns show the sizes of the operands issued to the Floating-Point Unit by the CPU. "D" indicates a 32-bit Double Word. "F" indicates that the instruction specifies an integer size for the operand (B = Byte, W = Word, D = Double Word). "F" indicates that the instruction specifies a Floating-Point size for the operand (F = 32-bit Standard Floating, L = 64-bit Long Floating).

The Returned Value Type and Destination column gives the size of any returned value and where the CPU places it. The PSR Bits Affected column indicates which PSR bits, if any, are updated from the Slave Processor Status Word (Figure 3-26).



Any operand indicated as being of type "F" will cause a transfer if the Register addressing mode specified. This is because the Floating-Point Register are physically on the Floating-Point Unit and therefore available without CPU assistance.

FIGURE 3-26. Slave Processor Status Word Format

TABLE 3-5. Floating-Point Instruction Protocols

Mnemonic	Operand 1 Class	Operand 2 Class	Operand 1 Issued	Operand 2 Issued	Returned Value Type and Dest	PSR Bits Affected
ADD	read, I	rmw, I	F	F	1 to Op. 2	none
SUB	read, I	rmw, I	F	F	1 to Op. 2	none
MUL	read, I	rmw, I	F	F	1 to Op. 2	none
DIV	read, I	rmw, I	F	F	1 to Op. 2	none
MOV	read, I	write, I	F	F	1 to Op. 2	none
MOV	read, I	write, I	F	F	1 to Op. 2	none
NEG	read, I	write, I	F	F	1 to Op. 2	none
CMV	read, I	read, I	F	F	NA	N, Z, L
FLOOR	read, I	write, I	F	F	1 to Op. 2	none
TRUNC	read, I	write, I	F	F	1 to Op. 2	none
ROUND	read, I	write, I	F	F	1 to Op. 2	none
MOV	read, F	write, L	F	F	1 to Op. 2	none
MOV	read, L	write, F	F	F	1 to Op. 2	none
MOV	read, I	write, I	F	F	1 to Op. 2	none
FSR	read, D	write, D	F	F	1 to Op. 2	none
FSR	read, D	write, D	F	F	1 to Op. 2	none
POLY	read, I	read, I	F	F	1 to Op. 2	none
DOT	read, I	read, I	F	F	1 to Op. 2	none
SCAL	read, I	read, I	F	F	1 to Op. 2	none
LOG	read, I	write, I	F	F	1 to Op. 2	none

Note:
D = Double Word
I = Integer size (B, W, D) specified in mnemonic.
F = Floating-Point type (F, L) specified in mnemonic.
NA = Not Applicable to this instruction.

3.0 Functional Description (Continued)

- 5) Set "Return Address" to the address of the first byte of the trapped instruction.
- 6) Perform Service (Vector, Return Address), Figure 3-24.
- 3.7.8.3 Trace Trap Sequence
 - 1) In the Processor Status Register (PSR), clear the P bit.
 - 2) Copy the PSR into a temporary register, then clear PSR bits S, U and T.
 - 3) Push the PSR copy onto the Interrupt Stack as a 16-bit value.
 - 4) "Vector" to 9.
 - 5) Set "Return Address" to the address of the next instruction.
 - 6) Perform Service (Vector, Return Address), Figure 3-24.

3.8 SLAVE PROCESSOR INSTRUCTIONS

The NS32EX16 supports only one group of instructions, the floating-point instruction set, as being executable by a slave processor. The floating-point instruction set is validated by the F bit in the CFG register.

If a floating-point instruction is encountered and the F bit in the CFG register is not set, a Trap (UND) will result, without any slave processor communication attempted by the CPU. This allows software emulation in case an external floating-point unit (FPU) is not used.

3.8.1 Slave Processor Protocol

Slave Processor instructions have a three-byte Basic Instruction field, consisting of an ID Byte followed by an Operation Word. The ID Byte has three functions:

- 1) It identifies the instruction as being a Slave Processor instruction.
- 2) It specifies which Slave Processor will execute it.
- 3) It determines the format of the following Operation Word of the instruction.

Upon receiving a Slave Processor instruction, the CPU initiates the sequence outlined in Figure 3-25. While applying Status Code 1111 (Broadcast ID, Section 3.4.1), the CPU transfers the ID Byte on the least-significant half of the Data Bus (AD0-AD7). All Slave Processors input this byte and decode it. The Slave Processor selected by the ID Byte is activated, and from this point the CPU is communicating only with it. If any other slave protocol was in progress (e.g., an aborted Slave instruction), this transfer cancels it.

Status Combinations:		
Send ID (ID):	Code 1111	
Xfer Operand (OP):	Code 1101	
Read Status (ST):	Code 1110	
Step	Status	Action
1	ID	CPU Sends ID Byte.
2	OP	CPU Sends Operation Word.
3	OP	CPU Sends Required Operands.
4	—	Slave Status Execution: CPU Prefetches.
5	—	Slave Puts SPC Low.
6	ST	CPU Reads Status Word (Trap? After Flag?)
7	OP	CPU Reads Result (If Any).

FIGURE 3-25. Slave Processor Protocol

The CPU next sends the Operation Word while applying Status Code 1101 (Transfer Slave Operand, Section 3.4.1). Upon receiving it, the Slave Processor decodes it, and at this point both the CPU and the Slave Processor are aware of the number of operands to be transferred and their sizes. The Operation Word is swapped on the Data Bus; that is, bits 0-7 appear on pins AD8-AD15 and 8-15 appear on pins AD0-AD7.

Using the Addressing Mode field within the Operation Word, the CPU starts fetching operands and issuing them to the Slave Processor. To do so, it references any Addressing Mode extensions which may be appended to the Slave Processor instruction. Since the CPU is solely responsible for memory accesses, these extensions are not sent to the Slave Processor. The Status Code applied is 1101 (Transfer Slave Processor Operand, Section 3.4.1).

After the CPU has issued the last operand, the Slave Processor starts the actual execution of the instruction. Upon completion, it will signal the CPU by pulsing SPC low.

While the Slave Processor is executing the instruction, the CPU is free to prefetch instructions into its queue. If it fills the queue before the Slave Processor finishes, the CPU will wait, applying Status Code 0011 (Waiting for Slave).

Upon receiving the pulse on SPC, the CPU uses SPC to read a Status Word from the Slave Processor, applying Status Code 1110 (Read Slave Status). This word has the format shown in Figure 3-26. If the O bit ("Out", Bit 0) is set, this indicates that an error was detected by the Slave Processor. The CPU will not continue the protocol, but will immediately trap through the Slave vector in the Interrupt Table. Certain Slave Processor instructions cause CPU PSR bits to be loaded from the Status Word.

The last step in the protocol is for the CPU to read a result, if any, and transfer it to the destination. The Read cycles from the Slave Processor are performed by the CPU while applying Status Code 1101 (Transfer Slave Operand).

3.0 Functional Description (Continued)

3.4.2 Basic Read and Write Cycles

The sequence of events occurring during a CPU access to either memory or peripheral device is shown in Figure 3-7 for a read cycle, and Figure 3-8 for a write cycle.

The cases shown assume that the selected memory or peripheral device is capable of communicating with the CPU at full speed. If not, then cycle extension may be requested through CWAIT and/or WAIT-2.

A full-speed bus cycle is performed in four cycles of the CTTL clock signal, labeled T1 through T4. Clock cycles not associated with a bus cycle are designated T1 (for "idle").

During T1, the CPU applies an address on pins AD0-AD15 and A16-A23. It also provides a low-going pulse on the ADS pin, which serves the dual purpose of informing external circuitry that a bus cycle is starting and of providing control to an external latch for demultiplexing Address bits 0-15 from the AD0-AD15 pins. See Figure 3-6.

However, using the ALE output signal is suggested for controlling the address latch. In normal CPU read cycles, and in external DMA cycles, ALE is asserted (high) at T1, and is deasserted (low) at T1 (see Figure 3-7). This eliminates the need for inverting the existing ADS off-chip to generate the address latch strobe, and removes the address latch strobe from the critical path to memory.

In CPU read and write cycles that access the on-chip FAM in slave cycles and in non-DMA idle states, ALE is always high. ALE is active (high) after reset. ALE is never Tri-State.

During this time also the status signals $\overline{D}(1)$, indicating the direction of the transfer, and $\overline{H}(1)$, indicating whether the high byte (AD8-AD15) is to be referenced, become valid.

During T2 the CPU switches the Data Bus, AD0-AD15, to either accept or present data. Note that the signals A16-A23 remain valid and need not be latched.

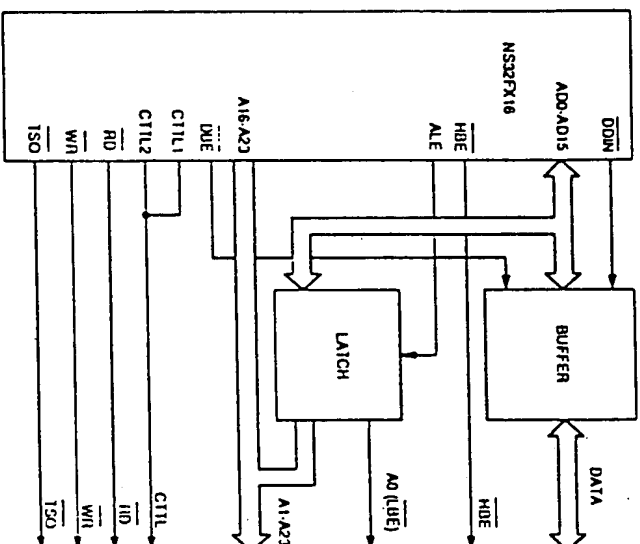


FIGURE 3-6. Bus Connections

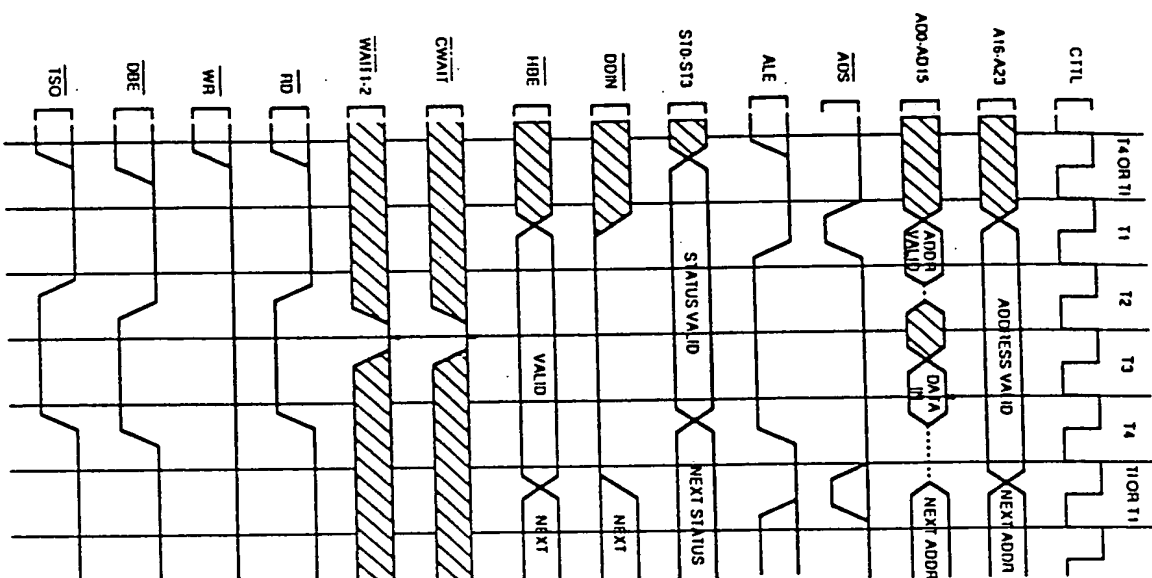


FIGURE 3-7. On-Chip Read Cycle Timing

3.7.7 Priority Among Exceptions

The NS32FX16 CPU internally prioritizes simultaneous interrupt and trap requests as follows:

- 1) Traps other than Trace (highest priority)

- 2) Non-Maskable Interrupt

- 3) Maskable Interrupts

- 4) Trace Trap (lowest priority)

3.7.8 Exception Acknowledge Sequences:

Detail Flow

For purposes of the following detailed discussion of interrupt and trap acknowledge sequences, a single sequence called "Service" is defined in Figure 3-24. Upon detecting any interrupt request or trap condition, the CPU first performs a sequence dependent upon the type of interrupt or trap. This sequence will include pushing the Processor Status Register and establishing a Vector and a Return Address. The CPU then performs the Service sequence.

3.7.8.1 Maskable/Non-Maskable Interrupt Sequence

This sequence is performed by the CPU when the NMI pin receives a falling edge, or the INTR pin becomes active with the PSRI bit set. The interrupt sequence begins either at the next instruction boundary or at the next interruptible point during its execution, as in the case of string or graphic instructions that have interior loops. The graphics instructions are interruptible.

1. If a String instruction was interrupted and not yet completed:

- a. Clear the Processor Status Register P bit.

- b. Set "Return Address" to the address of the first byte of the interrupted instruction. Otherwise, set "Return Address" to the address of the next instruction.

2. Copy the Processor Status Register (PSR) into a temporary register, then clear PSR bits S, U, T, P and I.

3. If the interrupt is Non-Maskable:

- a. Read a byte from address FFFFD016, applying Status Code 0100 (Interrupt Acknowledge, Master; Section 3.4.1). Discard the byte read.

- b. Set "Vector" to 1.

- c. Go to Step 6.

4. If the interrupt is Non-Vectored:

- a. Read a byte from address FFFFD016, applying Status Code 0100 (Interrupt Acknowledge, Master; Section 3.4.1). Discard the byte read.

- b. Set "Vector" to 0.

- c. Go to Step 6.

5. Here the interrupt is Vectored. Read "Byte" from address FFFFD016, applying Status Code 0100 (Interrupt Acknowledge, Master; Section 3.4.1).
6. If "Byte" = 0, then set "Vector" to "Byte" and go to Step 8.

7. If "Byte" is in the range -16 through -1, then the interrupt source is cascaded. (More negative values are reserved for future use.) Perform the following:

- a. Read the 32-bit Cascade Address from memory. The address is calculated as $\text{NTBASE} \times 4 + \text{Byte}$.

- b. Read "Vector" applying the Cascade Address just read and Status Code 0101 (Interrupt Acknowledge, Cascaded; Section 3.4.1).

8. Push the PSR copy (from Step 2) onto the interrupt Stack as a 16-bit value.

9. Perform Service (Vector, Return Address).

Figure 3-24.

Service (Vector, Return Address):

- 1) Read the 32-bit External Procedure Descriptor for the Interrupt Dispatch Table: address is Vector * 4, INTBASE Register contents.

- 2) Move the Module field of the Descriptor into the temporary MOD Register.

- 3) Read the Program Base pointer from memory address MOD * 8, and add it to the Offset field from the Descriptor, placing the result in the Program Counter.

- 4) Read the new Static Base pointer from the memory address contained in MOD, placing it into the SB Register.

- 5) Flush Queue: Non-sequentially fetch first instruction of interrupt routine.

- 6) Push MOD Register onto the interrupt Stack as a 16-bit value. (The PSR has already been pushed as a 16-bit value.)

- 7) Push the Return Address onto the interrupt Stack as a 32-bit quantity.

- 8) Copy temporary MOD Register to MOD Register.

FIGURE 3-24: Service Sequence involved during All Interrupt/Trap Sequences

3.7.8.2 Trap Sequence: Traps Other Than Trace

- 1) Restore the currently selected Stack Pointer and the Processor Status Register to their original values at the start of the trapped instruction.

- 2) Set "Vector" to the value corresponding to the trap type.

- SLAVE: Vector = 3
ILL: Vector = 4
SVC: Vector = 5
DIVZ: Vector = 6
FLO: Vector = 7
BPT: Vector = 8
LMD: Vector = 10.

- 3) Copy the Processor Status Register (PSR) into a temporary register, then clear PSR bits S, U, P and T.

- 4) Push the PSR copy onto the interrupt Stack as a 16-bit value.

3.0 Functional Description (Continued)

3.7.4 Non-Maskable Interrupt

The Non-Maskable Interrupt is triggered whenever a falling edge is detected on the $\overline{\text{NMI}}$ pin. The CPU performs an Interrupt Acknowledge. Master bus cycle when processing of this interrupt actually begins. The Interrupt Acknowledge cycle differs from that provided for Maskable Interrupts in that the address presented is FFF0016. The vector value used for the Non-Maskable Interrupt is taken as 1, regardless of the value read from the bus.

The service procedure returns from the Non-Maskable Interrupt using the Return from Trap (RETT) instruction. No special bus cycles occur on return.

For a full sequence of events in processing the Non-Maskable Interrupt, see Section 3.7.8.1.

3.7.5 Traps

Traps are processing exceptions that are generated as direct results of the execution of an instruction. The Return Address pushed by any trap except Trap (TRC) is the address of the first byte of the instruction during which the trap occurred. Traps do not disable interrupts, as they are not associated with external events. Traps recognized by NS32FX16 CPU are:

Trap (SLAVE): An exceptional condition was detected by the Floating Point Unit during the execution of a Slave Instruction. This trap is requested via the Status Word returned as part of the Slave Processor Protocol (Section 3.8.1).

Trap (ILL): Illegal operation. A privileged operation was attempted while the CPU was in User Mode (PSR bit U=1).

Trap (SVC): The Supervisor Call (SVC) instruction was executed.

Trap (OVZ): An attempt was made to divide an integer by zero. (The SLAVE trap is used for Floating Point division by zero.)

Trap (FLAG): The FLAG instruction detected a "1" in the PSR F bit.

Trap (BPT): The Breakpoint (BPT) instruction was executed.

Trap (TRC): The instruction just completed is being traced. See Section 3.7.6.

Trap (UND): An undefined opcode was encountered by the CPU.

3.7.6 Instruction Tracing

Instruction tracing is a feature that can be used during debugging to single-step through selected portions of a program. Tracing is enabled by setting the T-bit in the PSR Register. When enabled, the CPU generates a Trace Trap (TRC) after the execution of each instruction.

At the beginning of each instruction, the T bit is copied into the PSR P (Trace Pending) bit. If the P bit is set at the end of an instruction, then the Trace Trap is activated. If any other trap or interrupt request is made during a traced instruction, its entire service procedure

is allowed to complete before the Trace Trap occurs. Each interrupt and trap sequence handles the P bit for proper tracing, guaranteeing only one Trace Trap per instruction, and guaranteeing that the Return Address pushed during a Trace Trap is always the address of the next instruction to be traced.

Due to the fact that some instructions can clear the T and P bits in the PSR, in some cases a Trace Trap may not occur at the end of the instruction. This happens when one of the privileged instructions, BICPSRW or LPIW PSR, is executed.

In other cases, it is still possible to guarantee that a Trace Trap occurs at the end of the instruction, provided that special care is taken before returning from the Trace Trap Service Procedure. In case a BICPSRW instruction has been executed, the service procedure should make sure that the T bit in the PSR copy saved on the interrupt stack is set before executing the RETT instruction to return to the program being traced. If the RETT or RETI instructions have to be traced, the Trace Trap Service Procedure should set the P and T bits in the PSR copy of the interrupt stack that is going to be restored in the execution of such instructions.

While debugging the NS32FX16 instructions which have interior loops (BBOR, BBXOR, BBAND, BBFOR, EXTLT, MOVAP, SBTPS, TBITS), special care must be taken with the single-step trap. If an interrupt occurs during a single-step of one of the graphics instructions, the interrupt will be serviced. Upon return from the interrupt service routine, the new NS32FX16 instruction will not be re-entered, due to a single-step trap. Both the $\overline{\text{NMI}}$ and $\overline{\text{INT}}$ interrupts will cause this behavior. Another single-step operation (S command in DEBUGMON16) will resume from where the instruction was interrupted. There are no side effects from this early termination, and the instruction will complete normally.

For all other Series 32000 instructions, a single-step operation will complete the entire instruction before trapping back to the debugger. On the instructions mentioned above, several single-step commands may be required to complete the instruction, ONLY when interrupts are occurring.

There are some suggested methods to give the appearance of single-stepping for these NS32FX16 instructions.

1. MON16 monitors the return from the single-step trap vector's PC value. If the PC has not changed since the last single-step command was issued, the single-step operation is repeated. It is also advisable to ensure that one of the NS32FX16 instructions is being single-stepped by inspecting the first byte of the address pointed to by the PC register. If it is 0x0E, then the instruction is an NS32FX16 specific instruction.

2. A breakpoint following the instruction would also trap after the instruction had completed.

Note: If instruction tracing is enabled while the next instruction is executed, the Trap (TRC) occurs after the next interrupt, when the interrupt service procedure has returned.

3.0 Functional Description (Continued)

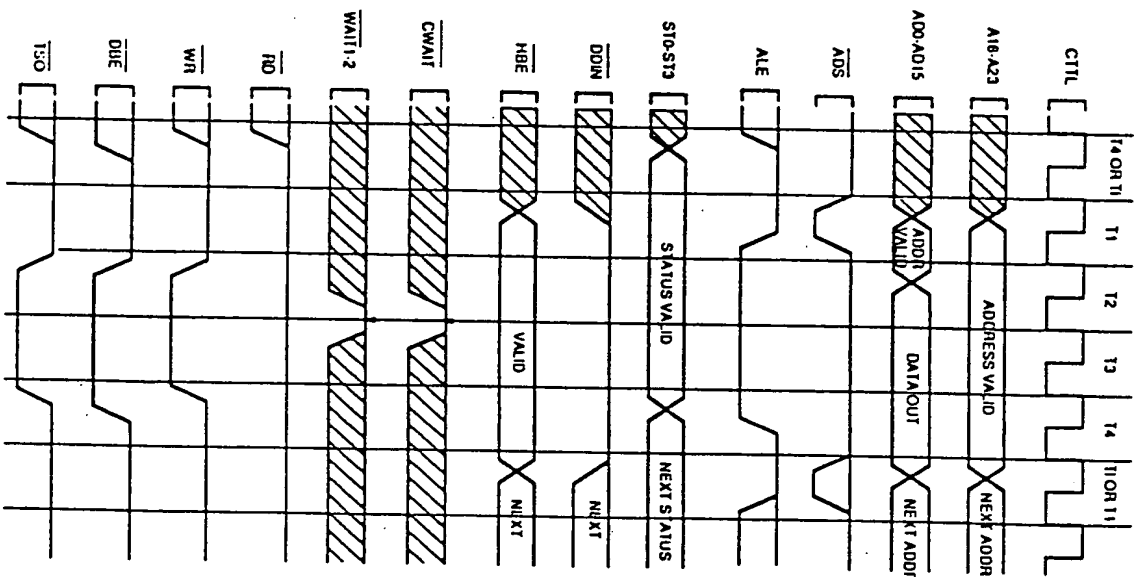


FIGURE 3-8. On-Chip Write Cycle Timing

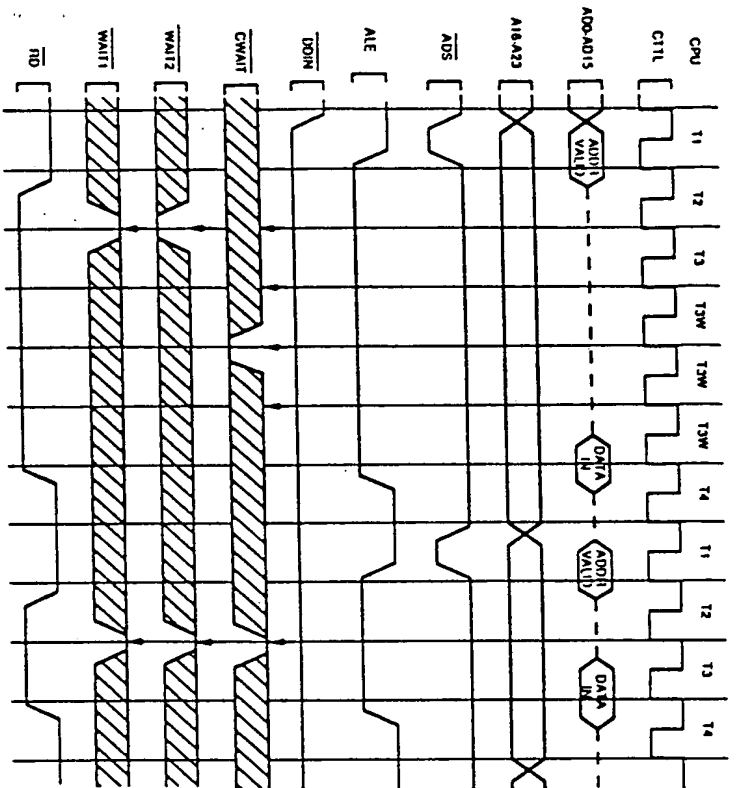


FIGURE 3-9. Extension of an OI-Chip Read Cycle

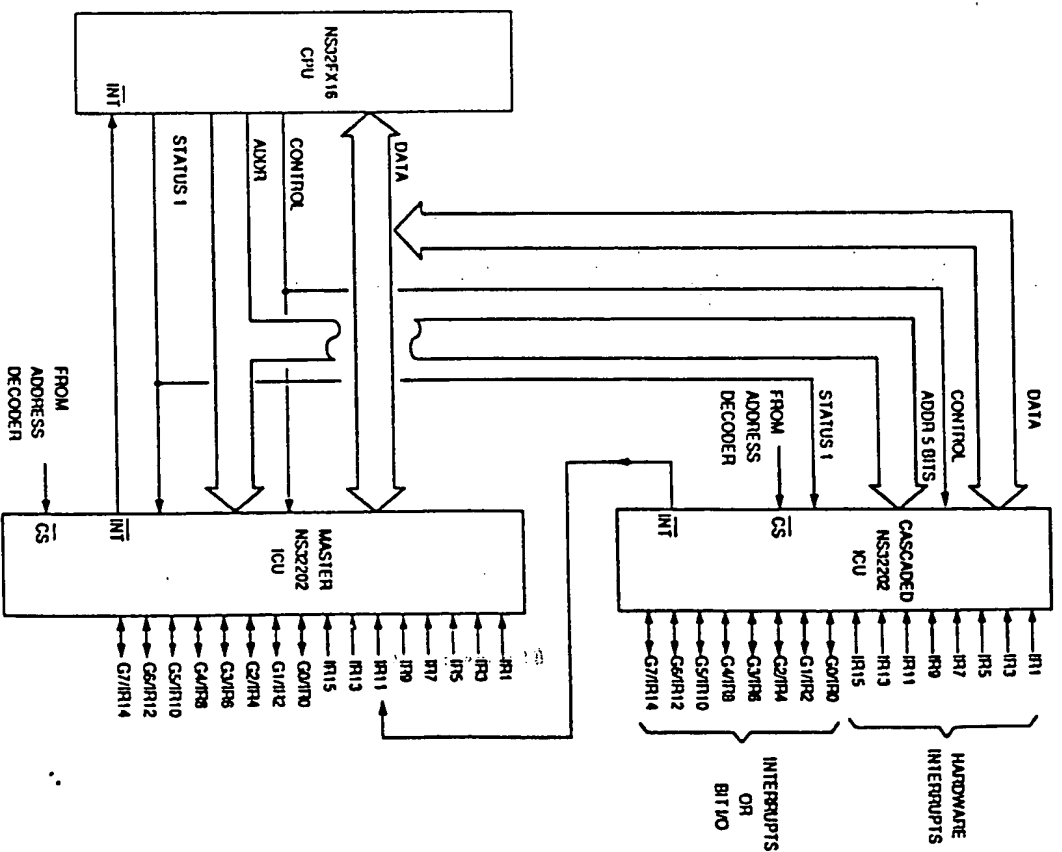


FIGURE 3-23. Cascaded Interrupt Control Unit Connections

In a system which uses cascading, two tasks must be performed upon initialization:

- 1) For each Cascaded KCU in the system, the Master KCU must be informed of the line number (0 to 15) on which it receives the cascaded requests.

2) A Cascade Table must be established in memory. The Cascade Table is located in a NEGATIVE direction from the location indicated by the CPU Interrupt Base (INTBASE) Register. Its entries are 32-bit addresses, pointing to the Vector Registers of each of up to 16 Cascaded KCUs.

Figure 3-18 illustrates the position of the Cascade Address Register. To find the Cascade Table entry for a Cascade Register, take its Master ICU line number (0 to 15) and subtract 16 from it, giving an index in the range -16 to 0. Multiply this value by 4, and add the resulting negative number to the contents of the INDBASE Register. The 32-bit entry at this address must be set to the address of the Hardware Vector Register of the Cascade ICU. This is referred to as the "Cascade Address".

Upon receipt of an interrupt request from a Cascaded ICU, the Master ICU interrupts the CPU and provides the negative Cascade Table index instead of a (positive) vector number. The CPU, seeing the negative value, uses it as an index into the Cascade Table and loads the **Cascade Address** from the referenced entry. Applying this address, the CPU performs an "interrupt" Acknowledge, Cascade" bus cycle, reading the final vector value. This vector is interpreted by the CPU as an unmasked byte and can therefore be in the range 0 through 255.

In returning from a Cascaded Interrupt, the service procedure executes the Return from Interrupt (IREI) instruction, as it would for any Maskable Interrupt. The CPU performs an "End of Interrupt, Master" bus cycle, whereupon the Master ICU again provides the negative Cascaded Table index. The CPU, seeing a negative Cascaded Table index, the corresponding Cascaded value, uses it to find the corresponding Cascaded address from the Cascade Table. Applying this address, it performs an "End of Interrupt, Cascaded" bus cycle, informing the Cascaded ICU of the completion of the service routine. The byte read from the Cascaded ICU is discarded.



FIGURE 3-22. Interrupt Control Unit Connections (16 Levels)

Note: If an interrupt must be masked off, the CPU can do so by setting the corresponding bit in the Interrupt Mask Register of the Interrupt Controller. However, if an interrupt is set pending during the CPU instruction that masks off that interrupt, the CPU may still perform an interrupt acknowledge cycle.

holding that instruction since it might have emptied the I/O line before the I/OU deserialized it. This could cause the I/OU to provide an invalid vector. To avoid this problem the above operation should be performed with the CPU interrupt disabled.

At this time the signals $\overline{\text{TSO}}$ (Timing State Output), $\overline{\text{DBE}}$ (Data Buffer Enable) and either $\overline{\text{RD}}$ (Read Strobe) or $\overline{\text{WR}}$ (Write Strobe) will also be activated.

The J3 state provides for access time requirements, and it occurs at least once in a bus cycle. At the end of T2, on the rising edge of CTTL, the CWait and Wait1-2 signals are sampled to determine whether the bus cycle will be extended. See Section 3.4.3.

If the CPU is performing a read cycle, the data bus (ADO-AD15) is sampled at the beginning of T4 on the rising edge of CTTL. Data must, however, be held a little longer to meet the data hold time requirements. The FUD signal is guaranteed not to go inactive before this time, so its rising edge can be safely used to disable the devices providing the input data.

The T4 state finishes the bus cycle. At the beginning of T4, the RD or WR, and I/O signals go inactive, and on the falling edge of CT11, D16 goes inactive, having provided for necessary data hold times. Data during write cycles remains valid from the CPU throughout T4. Note that the Bus Status lines (SIO-S13) change at the beginning of T4, anticipating the following bus cycle (if any).

To allow sufficient access time for any speed of memory or peripheral device, the NS32FX16 provides for extension of a bus cycle. Any type of bus cycle except a Slave Processor cycle can be extended.

In *Figures 3-7* and *3-8*, note that during T3 all bus control signals from the CPU are flat. Therefore, a bus cycle can be clearly extended by causing the T3 state to be repeated. This is the purpose of the WAIT-2 and WAIT input signals.

At the end of state T2, on the rising edge of CTL, WAIT1.2 and WAIT are sampled.

If any of these signals are active, the bus cycle will be extended by at least one clock cycle. Thus, one or more additional T2 states (also called wait state T3W) will be inserted after the next T1 state. Any combination of the above signals can be activated at one time. However, the WAIT1:2 inputs are only sampled by the CPU at the end of state T2. They are ignored at all other times.

The WAIT1-2 inputs are binary weighted, and can be used to insert up to 3 wait states, according to the following table.

WAIT2	WAIT1	Number of Wall States
HIGH	HIGH	0
HIGH	LOW	1
LOW	HIGH	2
LOW	LOW	3

CWAIT causes wait states to be inserted continuously as long as it is sampled active. It is normally used when the number of wait states to be inserted in the CPU bus cycle is not known in advance.

The following sequence shows the CPU response to the WAIT1-2 and CWAIT inputs.

1. Start bus cycle.
2. Sample $\overline{G}WAI_{1-2}$ and $\overline{G}WAI_{AT}$ at the end of state 12.
3. If the $\overline{G}WAI_{1-2}$ inputs are both inactive, then go to step 6.
4. Insert the number of wait states selected by \overline{WAI}_{1-2} .
5. Sample $\overline{G}WAI_{AT}$ again.
6. If $\overline{G}WAI_{AT}$ is not active, then go to step 8.
7. Insert one wait state and then go to step 5.
8. Complete bus cycle.

Figure 3-9 shows a bus cycle extended by three wait states, two of which are due to WAIT2 and one of which is due to WAIT.

The 24-bit address provided by the NS32CFX16 is a byte address; that is, it uniquely identifies one of up to $16,777,216$ eight-bit memory locations. An important feature of the NS32CFX16 is that the presence of a 16-bit data bus imposes no restrictions on data alignment; any data item, regardless of size, may be placed starting at any memory address. The NS32CFX16 provides a special control signal, High Byte Enable (HtE), which facilitates individual byte addressing on a 16-bit bus.

Memory is organized as two eight-bit banks, each bank receiving the word address (A1-A20) in parallel. One address, connected to Data Bus pins AD0-AD7, is enabled to respond to even byte addresses; i.e., when the least significant address bit (A0) is low. The other bank, connected to Data Bus pins AD8-AD15, is enabled when A0 is high. See *Figure 3-10*.

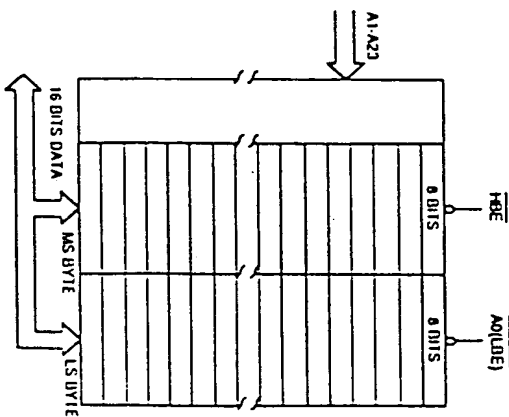


FIGURE 3-10. Memory Interface

3.0 FUNCTIONAL DESCRIPTION (Continued)

Any bus cycle falls into one of three categories: Even Byte Access, Odd Byte Access, and Even Word Access. All accesses to any data type are made up of sequences of these cycles. Table 3-2 gives the state of A0 and H0E for each category.

TABLE 3-2. Bus Cycle Categories

Category	H0E	A0
Even Byte	1	0
Odd Byte	0	1
Even Word	0	0

Accesses of operands requiring more than one bus cycle are performed sequentially, with no idle T-States separating them. The number of bus cycles required to transfer an operand depends on its size and its alignment (i.e., whether it starts on an even byte address or an odd byte address). Table 3-3 lists the bus cycle performed for each situation. For the timing of A0 and H0E, see Section 3.4.2.

3.4.4.1 Bit Accesses

The Bit Instructions perform byte accesses to the byte containing the designated bit. The Test and Set Bit instruction (SBIT), for example, reads a byte, alters it, and rewrites it, having changed the contents of one bit.

3.4.4.2 Bit Field Accesses

An access to a Bit Field in memory always generates a Double-Word transfer at the address containing the least significant bit of the field. The Double Word is read by an Extract instruction; an Insert instruction reads a Double Word, modifies it, and rewrites it.

3.4.4.3 Extended Multiple Accesses

The Multiply Extended Integer (MEI) instruction will return a result which is twice the size in bytes of the operand it reads. If the operand is in memory, the most significant half of the result is written first (at the higher address), then the least significant half.

3.4.5 Instruction Fetches

Instructions for the NS32FX16 CPU are "prefetched"; that is, they are input before being needed into the next available entry of the eight-byte Instruction Queue. The CPU performs two types of Instruction Fetch cycles: Sequential and Non-Sequential. These can be distinguished from each other by their differing status combinations on pins ST0-ST3 (Section 3.4.1).

A Sequential Fetch will be performed by the CPU whenever the Data Bus would otherwise be idle and the Instruction Queue is not currently full. Sequential Fetches are always Even Word Read cycles (Table 3-2). A Non-Sequential Fetch occurs as a result of any break in the normally sequential flow of a program. Any jump or branch instruction, a trap or an interrupt, will cause the next Instruction Fetch cycle to be Non-Sequential.

In addition, certain instructions flush the instruction queue, causing the next instruction fetch to display Non-Sequential status. Only the first bus cycle after a break displays Non-Sequential status, and that cycle is either an Even Word Read or an Odd Byte Read, depending on whether the destination address is even or odd.

3.4.6 Interrupt Control Cycles

Activating the \overline{INT} or \overline{NMI} pin on the CPU will initiate one or more bus cycles whose purpose is interrupt control rather than the transfer of instructions or data. Execution of the Return from Interrupt instruction (RETI) will also cause Interrupt Control bus cycles. These differ from instruction or data transfers only in the status presented on pins ST0-ST3. All Interrupt Control cycles are single-byte Read cycles.

Table 3-4 shows the Interrupt Control sequences associated with each interrupt and with the return from its service routine. For full details of the NS32FX16 Interrupt structure, see Section 3.7.

END OF INTERRUPT
BUS CYCLE

INTERRUPT CONTROL UNIT

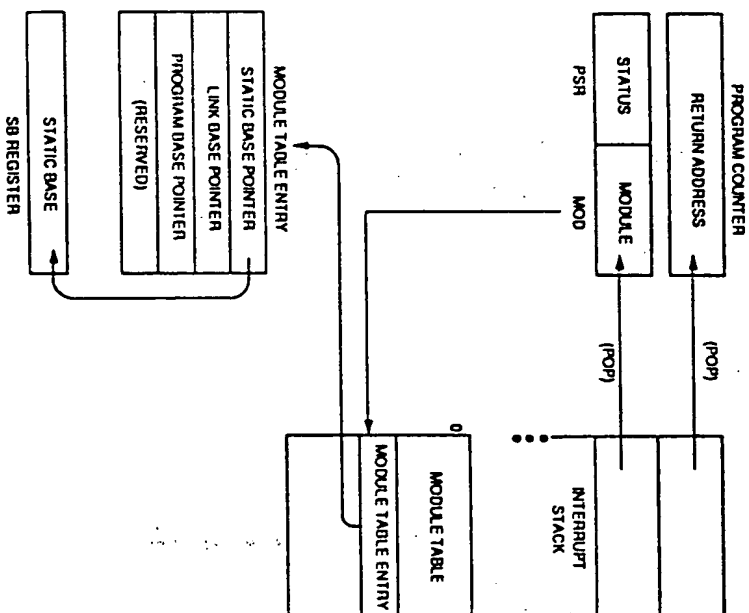


FIGURE 3-21. Return from Interrupt (RETI) Instruction Flow

3.7.3.2 Vectored Mode: Non-Cascaded Case

In the Vectored mode, the CPU uses an Interrupt Control Unit (ICU) to prioritize up to 16 interrupt requests. Upon receipt of an interrupt request on the \overline{INT} pin, the CPU performs an "Interrupt Acknowledge Master" bus cycle reading a vector value from the low-order byte of the Data Bus. This vector is then used as an index into the Dispatch Table in order to find the External Procedure Descriptor for the proper interrupt service procedure. The service procedure eventually returns via the Return from Interrupt (RETI) instruction, which performs an End of Interrupt bus cycle, informing the ICU that it may re-prioritize any interrupt requests still pending. The ICU provides the vector number again, which the CPU uses to determine whether it needs also to inform a Cascaded ICU.

In a system with only one ICU (16 levels of interrupt), the vectors provided must be in the range 0 through 127; that is, they must be positive numbers in eight bits. By providing a negative vector number, an ICU flags the interrupt source as being a Cascaded ICU (see below).

3.7.3.3 Vectored Mode: Cascaded Case

In order to allow up to 256 levels of interrupt, provision is made both in the CPU and in the NS3202 Interrupt Control Unit (ICU) to transparently support cascading. Figure 3-23 shows a typical cascaded configuration. Note that the interrupt output from a Cascaded ICU goes to an Interrupt Request input of the Master ICU, which is the only ICU which drives the CPU \overline{INT} pin.

3.0 Functional Description (Continued)

3.7.2 Returning from an Exception Service Procedure

To return control to an interrupted program, one of two instructions can be used: RETT (Return from Trap) and RETI (Return from Interrupt).

RETT is used to return from any trap or a non-maskable interrupt service procedure. Since some traps are often used deliberately as a call mechanism for supervisor mode procedures, RETT can also adjust the Stack Pointer (SP) to discard a specified number of bytes from the original stack as surplus parameter space.

RETI is used to return from a maskable interrupt service procedure. A difference of RETI, RETI also informs any external interrupt control units that interrupt service has completed. Since interrupts are generally asynchronous external events, RETI does not discard parameters from the stack.

Both of the above instructions always restore the PSR, MOD, PC and SB registers to their previous contents.

3.7.3 Maskable Interrupts

The INT pin is a level-sensitive input. A continuous low level is allowed for generating multiple interrupt requests. The input is maskable, and is therefore enabled to generate interrupt requests only while the Processor Status Register bit is set. The 1 bit is automatically cleared during service of an INT or NMI request, and is restored to its original setting upon return from the interrupt service routine via the RETT or RETI instruction.

The INT pin may be configured via the SETCFG instruction as either Non-Vectored (CFG Register bit I=0) or Vectored (bit I=1).

3.7.3.1 Non-Vectored Mode

In the Non-Vectored mode, an interrupt request on the INT pin will cause an Interrupt Acknowledge bus cycle, but the CPU will ignore any value read from the bus and use instead a default vector of zero. This mode is useful for small systems in which hardware interrupt prioritization is unnecessary.

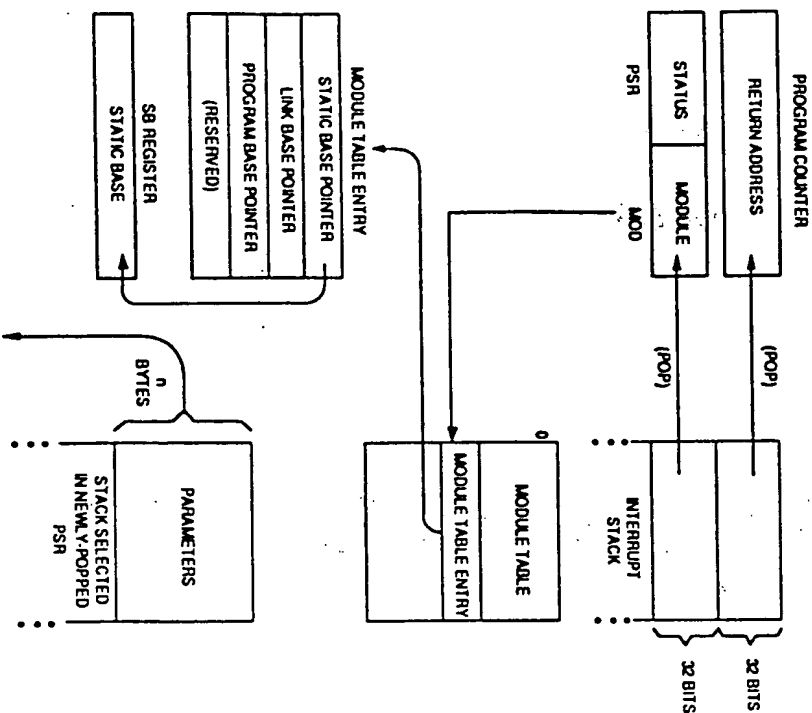


FIGURE 3.70. Return from Trap (RETT) Instruction Flow

3.0 Functional Description (Continued)

TABLE 3.3. Access Sequences

Cycle Type	Address	HIDE	A0	High Bus	Low Bus
A. Odd Word Access Sequence					
1	Odd Byte	A	0	1	Byte 0 Don't Care
2	Even Byte	A+1	1	0	Don't Care Byte 1
B. Even Double-Word Access Sequence					
1	Even Word	A	0	0	Byte 0 Byte 2
2	Even Word	A+2	0	0	Byte 1 Byte 3
C. Odd Double-Word Access Sequence					
1	Odd Byte	A	0	1	Byte 0 Don't Care
2	Even Word	A+1	0	0	Byte 2 Byte 1
3	Even Byte	A+3	1	0	Don't Care Byte 3
D. Even Quad-Word Access Sequence					
1	Even Word	A	0	0	Byte 0 Byte 2
2	Even Word	A+2	0	0	Byte 1 Byte 3
3	Even Word	A+4	0	0	Byte 4 Byte 6
4	Even Word	A+6	0	0	Byte 5 Byte 7
E. Odd Quad-Word Access Sequence					
1	Odd Byte	A	0	1	Byte 0 Don't Care
2	Even Word	A+1	0	0	Byte 2 Byte 1
3	Even Byte	A+3	1	0	Don't Care Byte 3
4	Odd Byte	A+4	0	1	Byte 4 Don't Care
5	Even Word	A+5	0	0	Byte 6 Byte 5
6	Even Byte	A+7	1	0	Don't Care Byte 7

Other bus cycles (instruction prefetch or slave) can occur here.

TABLE 3-4. Interrupt Sequences

Cycle	Status	Address	DM	IDE	A0	High Bus	Low Bus
A. Non-Maskable Interrupt Control Sequence							
Interrupt Acknowledge							
1	0100	FFFF0016	0	1	0	Don't Care	Don't Care
Interrupt Return							
None; Performed through Return from Trap (RETT) instruction.							
B. Non-Vectored Interrupt Control Sequence							
Interrupt Acknowledge							
1	0100	FFFF0016	0	1	0	Don't Care	Don't Care
Interrupt Return							
None; Performed through Return from Trap (RETT) instruction.							
C. Vectored Interrupt Sequence: Non-Cascaded							
Interrupt Acknowledge							
1	0100	FFFF0016	0	1	0	Don't Care	Vector: Range: 0-127
Interrupt Return							
1	0110	FFFF0016	0	1	0	Don't Care	Vector: Same as in Previous Int. Ack. Cycle
D. Vectored Interrupt Sequence: Cascaded							
Interrupt Acknowledge							
1	0100	FFFF0016	0	1	0	Don't Care	Cascade Index: range - 16 to -1
2	0101	Cascade Address	0	1 ^a	0 ^a	Vector range 0-255; on appropriate half of Data Bus for even/odd address	
Interrupt Return							
1	0110	FFFF0016	0	1	0	Don't Care	Cascade Index: same as in previous Int. Ack. Cycle

(The CPU here uses the Cascade Index to find the Cascade Address.)

2 0111 Cascade Address 0^a 1^a 0^a Don't Care Don't Care

^a If the Cascaded ICU address is Even (A0 is low), then the CPU applies \overline{HIDE} high and reads the vector number from bit 0-7 of the Data Bus. If the address is Odd (A0 is high), then the CPU applies \overline{HIDE} low and reads the vector number from bit 8-15 of the Data Bus. The vector number may be in the range 0-255.

This process is illustrated in Figure 3-19 from the view point of the programmer.

Details on the sequences of events in processing interrupts and traps are given in the following sections.

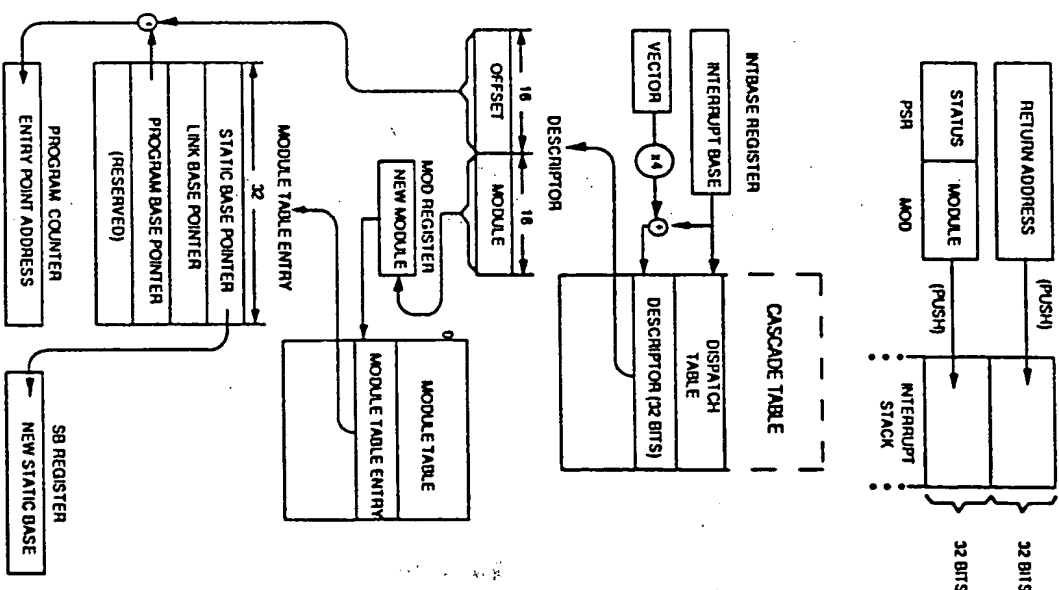


FIGURE 3-19. Exception Acknowledge Sequence

3.0 Functional Description (Continued)

3.7 EXCEPTION PROCESSING

Exceptions are special events that alter the sequence of instruction execution. The CPU recognizes two basic types of exceptions: interrupts and traps.

An interrupt occurs in response to an event signaled by activating the NMI or INT input signals. Interrupts are typically requested by peripheral devices that require the CPU's attention.

Traps occur as a result of exceptional conditions (e.g., attempted division by zero) or of specific instructions whose purpose is to cause a trap to occur (supervisor call instruction).

When an exception is recognized, the CPU saves the PC, PSR and the MOD register contents on the interrupt stack, and then it transfers control to an exception service procedure.

Details on the operations performed in the various cases by the CPU to enter and exit the exception service procedure are given in the following sections.

It is to be noted that the reset operation is not treated here as an exception, even though, like any exception, it alters the instruction execution sequence. This is because the CPU handles reset in a significantly different way than it handles exceptions.

Refer to Section 3.3 for details on the reset operation.

3.7.1 Exception Acknowledge Sequence

When an exception is recognized, the CPU goes through three major steps:

1) Adjustment of Registers.

Depending on the source of the exception, the CPU may restore and/or adjust the contents of the Program Counter (PC), the Processor Status Register (PSR) and the currently-selected Stack Pointer (SP). A copy of the PSR is made and the PSR is then set to reflect Supervisor Mode and solution of the interrupt stack.

2) Vector Acquisition

A Vector is either obtained from the Data Bus or is supplied by default.

3) Service Call

The Vector is used as an index into the Interrupt Dispatch Table, whose base address is taken from the CPU Interrupt Base (INTBASE) Register. See Figure 3-18. A 32-bit External Procedure Call is performed using it. The MOD Register (16 bits) and Program Counter (32 bits) are pushed on the Interrupt Stack.

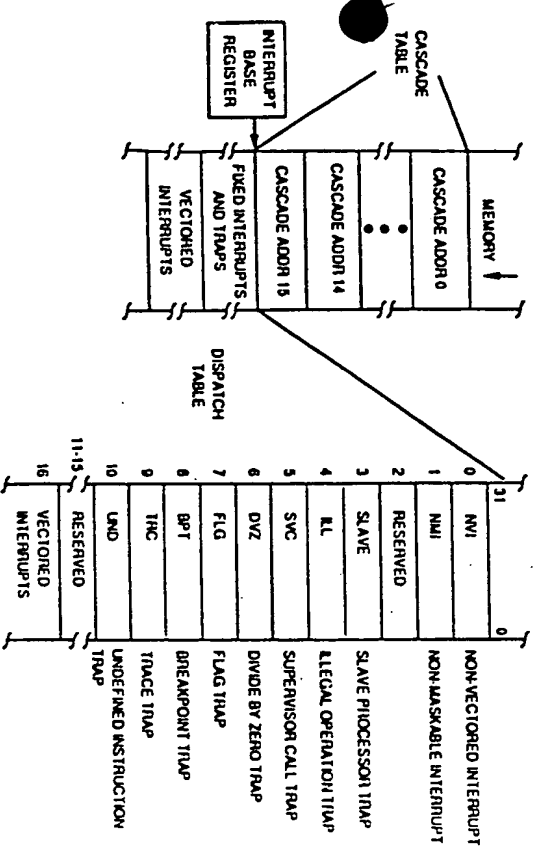


FIGURE 3-18. Interrupt Dispatch and Cascade Tables

3.0 Functional Description (Continued)

3.4.7 On-Chip Bus Cycles

The bus cycles accessing registers of the on-chip FAX Accelerator Module do not involve any off-chip resource. However, for observability reasons, the NS32FX16's bus interface provides all the necessary information in order to allow a debug or trace device (e.g., ISE) to track an on-chip bus transaction.

An on-chip bus transaction is very similar (hence) to an off-chip bus transaction. However, the ADS, AD, WI, TSO, and DDE outputs are not asserted by the CPU. Instead, the NS32FX16 asserts a special output, IAS.

During write cycles to on-chip addresses, the data to be written can be observed on ADO-AD15.

Access to the FAM registers while it is executing a vector operation are delayed (as if the CWMIT input is active). When the FAM finishes the operation, access to the registers proceeds. Those wait states cannot be observed on external pins.

The address on ADO-AD15 and A16-A23 during internal reference is the 24 least significant bits of the addressed internal register address.

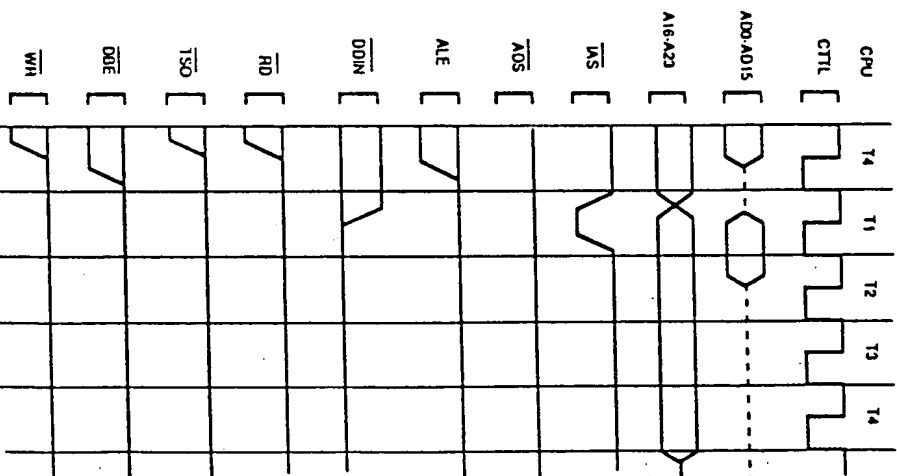


FIGURE 3-11 (a). On-Chip Read Cycle

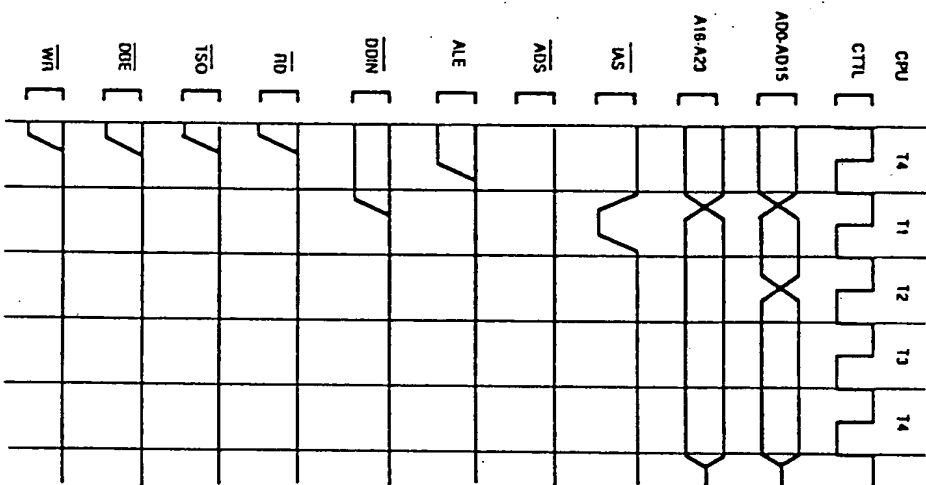


FIGURE 3-11 (b). On-Chip Write Cycle

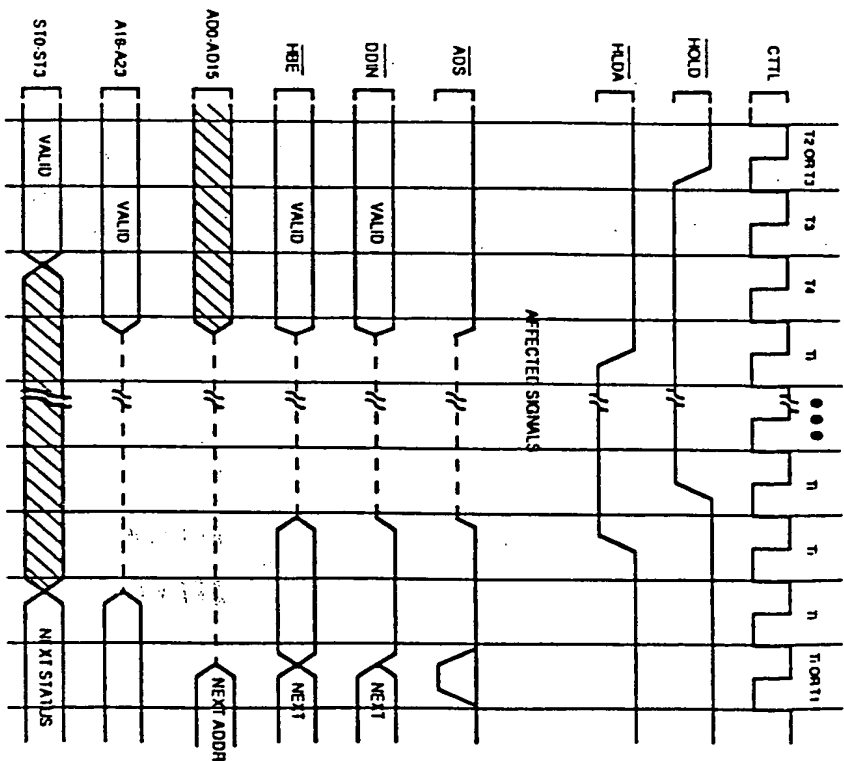


FIGURE 3-17. HOLD Timing, Bus Initially Not Idle

3.6 INSTRUCTION EXECUTION AND STATUS

In addition to the four bits of Bus Cycle status (ST0-ST3), the NS32FX16 CPU also presents Instruction Status information on three separate pins. These pins differ from ST0-ST3 in that they are synchronous to the CPU's internal instruction execution section rather than to its bus interface section.

PFS (Program Flow Status) is pulsed low as each instruction begins execution. It is intended for debugging purposes.

UIS originates from the U bit of the Processor Status Register, and indicates whether the CPU is currently running in User or Supervisor mode. Although it is not synchronous to bus cycles, there are guarantees on its

validity during any given bus cycle. See the Timing Specifications in Section 4.

IC0 (Interlocked Operation) is activated during an SBIT (Set Bit, Interlocked) or CBIT (Clear Bit, Interlocked) instruction. It is made available to external bus arbitration circuitry in order to allow these instructions to implement the semaphore primitive operations for multi-processor communication and resource sharing. IC0 is guaranteed to be active during the operand accesses performed by the interlocked instructions.

Note: The acknowledge of IC0 is on a cycle by cycle basis. Therefore, it is possible to have IC0 active when an interlocked operation is in progress. In this case, IC0 remains low and the interlocked instruction continues only after IC0 is de-asserted.

3.0 Functional Description (Continued)

(HOLD Acknowledge) pins. By asserting $\overline{\text{HOLD}}$ low, an external device requests access to the bus. On receipt of $\overline{\text{HOLD}}$ from the CPU, the device may perform bus cycles, as the CPU at this point has set A00-AD15 , A16-A23 and H0E to the TRI-STATE condition and has switched ADS and DDIN to the input mode. The CPU now monitors ADS and DDIN from the external device to generate the relevant strobe signals (i.e., TSO , DBE , HD or WH). To return control of the bus to the CPU, the device sets $\overline{\text{HOLD}}$ inactive, and the CPU acknowledges the return of the bus by setting $\overline{\text{H0LD}}$ inactive.

How quickly the CPU releases the bus depends on whether it is idle on the bus at the time the $\overline{\text{HOLD}}$ request is made, as the CPU must always complete the current bus cycle. Figure 3-16 shows the timing sequence when the CPU is idle. In this case, the CPU grants the bus during the immediately following clock

cycle. Figure 3-17 shows the sequence if the CPU is using the bus at the time that the $\overline{\text{HOLD}}$ request is made. If the request is made during or before the clock cycle shown (two clock cycles before T4), the CPU will release the bus during the clock cycle following T4 . If the request occurs closer to T4 , the CPU may already have decided to initiate another bus cycle. In that case, it will not grant the bus until after the next T4 state. Note that this situation will also occur if the CPU is idle on the bus but has initiated a bus cycle internally.

Note 1: During DMA cycles the $\overline{\text{WAIT}}/2$ signals should be kept inactive, unless they are also monitored by the DMA controller. If wait states are required, $\overline{\text{CWAIT}}$ should be used.

Note 2: The logic value of the status pins, ST0-ST3 , is undefined during DMA activity.

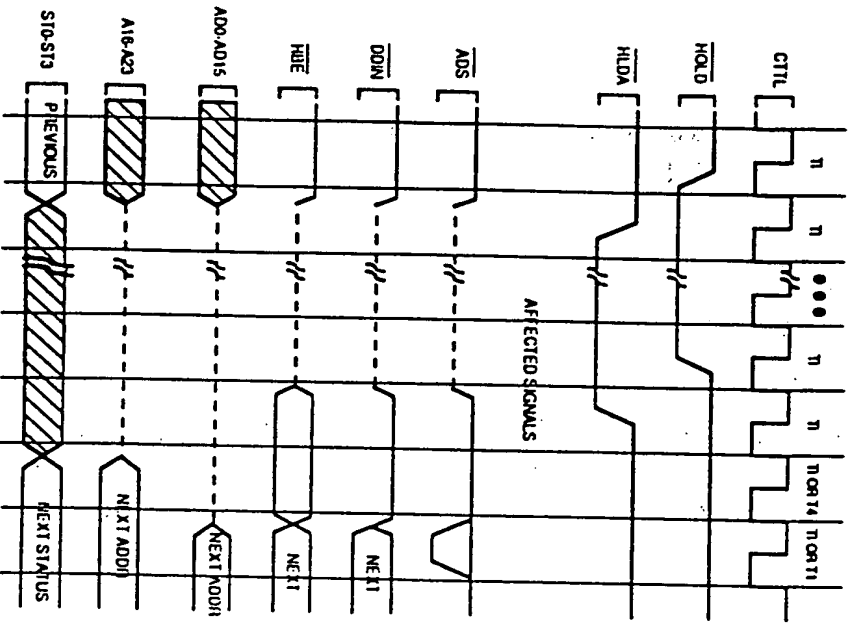


FIGURE 3-16. HOLD Timing, Bus Initially Idle

3.0 Functional Description (Continued)

3.4.8 Initiated by On-Chip DMA Controller
On-chip DMA Controller requests are handled by the bus access control mechanism. When granted with the bus (if it is asserted), the DMA Controller issues ADS and DDIN signals to the CPU and drives the address/

data buses. The CPU supports the DMA bus cycle by generating bus control signals as HD , WH , TSO and DBE . As a result, the DMA bus cycles are very similar to the CPU bus cycles. This simplifies the memory system design significantly.

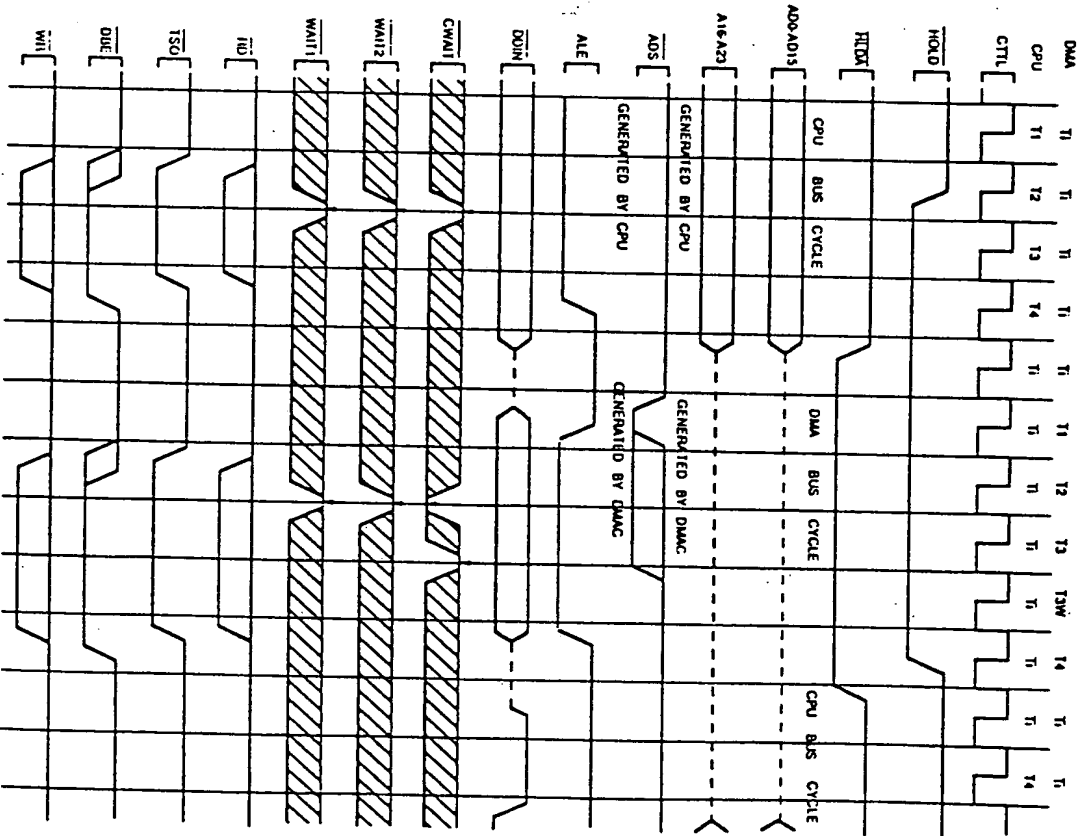


FIGURE 3-12. DMAC Initiated Bus Cycle

3.0 Functional Description (Continued)

3.4.9 Slave Processor Communication

The \overline{SPC} pin is used as the data strobe for Slave Processor transfers. In a Slave Processor bus cycle, data is transferred on the Data Bus ($AD0-AD15$), and the status lines $S10-S13$ are monitored by the Slave Processor.

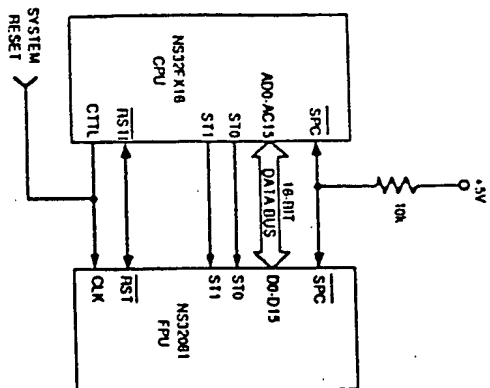


FIGURE 3-13(e). System Connection Diagram with the NS32081 FPU

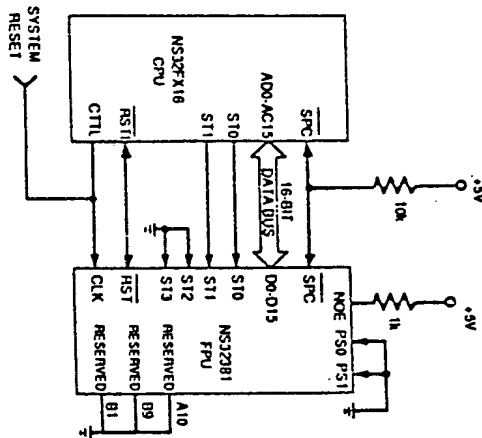
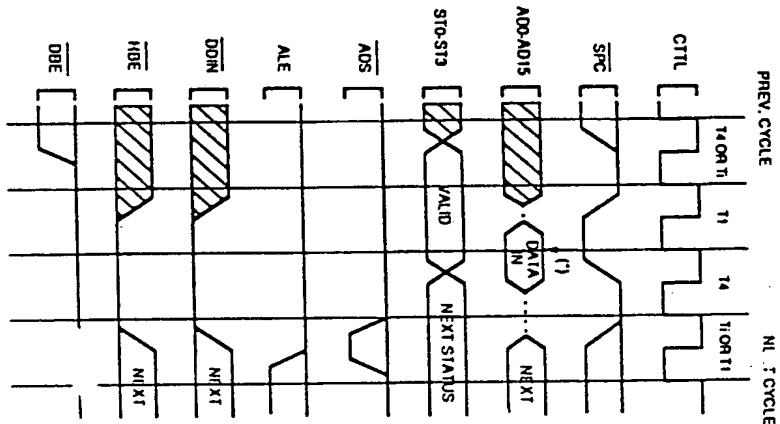


FIGURE 3-13(b). System Connection Diagram with the NS32381 FPU

In order to determine the type of transfer being performed, \overline{SPC} is bidirectional, but is driven by the CPU during all Slave Processor bus cycles. See Section 3.8 for full protocol sequences.



(1): CPU SAMPLES DATA BUS HERE
FIGURE 3-14. Slave Processor Read Cycle

3.4.9.1 Slave Processor Bus Cycles

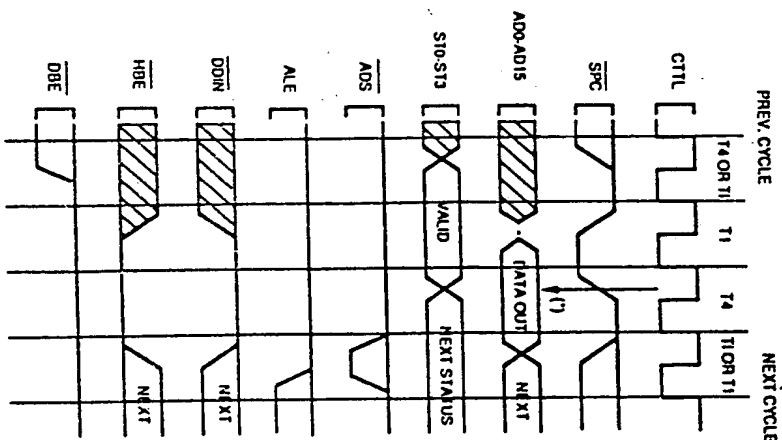
A Slave Processor bus cycle always takes exactly two clock cycles, labeled T1 and T4 (see Figures 3-14 and 3-15). During a Read cycle, \overline{SPC} is active from the beginning of T1 to the beginning of T4, and the data is sampled at the end of T1. The Cycle Status pins lead the cycle by one clock period, and are sampled at the leading edge of \overline{SPC} . During a Write cycle, the CPU applies data and activates \overline{SPC} at T1, removing \overline{SPC} at T4. The Slave Processor latches status on the leading edge of \overline{SPC} and latches data on the trailing edge. The CPU does not pulse the Address Strobe (\overline{ADS}) and no bus signals are generated. The ALE signal remains high during the slave cycle. The direction of a transfer is determined by the sequence ("protocol") established by the instruction under execution, but the CPU indicates the direction on the \overline{DDIN} pin for hardware debugging purposes.

3.4.9.2 Slave Operand Transfer Sequence

A Slave Processor operand is transferred in one or more Slave bus cycles. A 16-bit operand is transferred on the least-significant byte of the Data Bus ($AD0-AD7$), and a Word operand is transferred on the entire bus. A Double Word is transferred in a consecutive pair of bus cycles, least-significant word first. A Quad Word is transferred in two pairs of Slave cycles, with other bus cycles possibly occurring between them. The word order is from least-significant word to most-significant.

3.5 BUS ACCESS CONTROL

The NS32EX16 CPU has the capability of relinquishing its access to the bus upon request from a DMA controller or another CPU. This capability is implemented on the \overline{HOLD} (Hold Request) and \overline{HILDA} (Hold Inhibit) pins.



(1): SLAVE PROCESSOR SAMPLES DATA BUS HERE
FIGURE 3-15. Slave Processor Write Cycle

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

☐ BLACK BORDERS

☒ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES

☐ FADED TEXT OR DRAWING

☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING

☐ SKEWED/SLANTED IMAGES

☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS

☐ GRAY SCALE DOCUMENTS

☐ LINES OR MARKS ON ORIGINAL DOCUMENT

☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY

☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.